

Compositional security and collateral leakage

N Bordenabe¹, A McIver¹, C Morgan², and T Rabehaja¹

¹ Dept. Computing, Macquarie University, Sydney

² Data61 & UNSW, Sydney

Abstract

In quantitative information flow we say that program Q is “at least as secure as” P just when the amount of secret information flowing from Q is never more than flows from P , with of course a suitable quantification of “flow”. This secure-refinement order \sqsubseteq is *compositional* just when $P \sqsubseteq Q$ implies $\mathcal{C}(P) \sqsubseteq \mathcal{C}(Q)$ for any context \mathcal{C} , again with a suitable definition of “context”.

Remarkable however is that leaks caused by executing P, Q might not be limited to their declared variables: they might impact correlated secrets in variables declared and initialised in some broader context to which P, Q do not refer even implicitly. We call such leaks *collateral* because their effect is felt in domains of which (the programmers of) P, Q might be wholly unaware: our inspiration is the “Dalenius” phenomenon for statistical databases [9, 11].

We show that a proper treatment of these collateral leaks is necessary for a compositional program semantics for read/write “open” programs. By adapting a recent Hidden-Markov denotational model for non-interference security [21, 22], so that it becomes “collateral aware”, we give techniques and examples (e.g. public-key encryption) to show how collateral leakage can be calculated and then bounded in its severity.

Keywords and phrases Quantitative information flow, program semantics, secure refinement.

1 Introduction

The problem of information disclosure in the context of statistical databases was formulated by Dalenius as an ideal privacy goal: “Nothing about an individual should be learnable from the database that cannot be learned without access to the database” [9]. Later he argued the infeasibility of such a strict goal; more recently Dwork [11] addressed the same concern demonstrating that whenever there is a (known) *correlation* between two pieces of information, anything learned about one piece implies that something might also be learned about the other. In secure programming generally, i.e. not only read-only databases, this corresponds to leaking information about a secret “high-level” variable X , which then consequentially leaks information about a different high-level variable Z that *does not appear in the program at all*, but is known via “auxiliary information” to be correlated with the initial value of X . Because of the generality of this programming-language perspective, we call this effect *collateral leakage*.

Here we study this broader phenomenon of collateral leakage in the general setting of read/write programs operating as “open systems”. Because [1] studied Dalenius leakage in terms of abstract channels (read-only programs), and [22] explored the semantics of information flow for read/write programs operating as “closed systems” without external correlation, this work can be seen as bringing the two together: we compute general bounds on collateral leakage for read/write open systems, and we adapt our earlier fully abstract program semantics to treat it compositionally.

The following “password” example helps to illustrate the issues.



licensed under Creative Commons License CC-BY

Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

// Password X is initially uniformly distributed over $\mathcal{X} = \{A, B, C\}$.	
“Lax” user	“Strict” user
$X \leftarrow [A, B, C]$ *	$X \leftarrow [X^+, X^-]$ §
$\text{leak } [X^+, X^-]$ †	$\text{leak } [X^+, X^-]$ †

* $[\dots]$ is the uniform distribution over $\{\dots\}$; and $X \leftarrow$ assigns to X from a distribution.
 § X^+ is the letter following X in \mathcal{X} (wrapping around), and X^- the preceding.
 † $\text{leak } [X^+, X^-]$ makes a (fair) choice secretly between X^+ or X^- , then emits the value somehow: it does not however indicate whether it chose X^+ or X^- to leak.
 Note that the X referred to in the leak statement is the updated value, after the $X \leftarrow$ update.

Lax may choose any new password, uniformly, including his current; but Strict must *change* his password, again uniformly. In both cases the distribution of the new X is again uniform: for Lax it is independent of X ’s initial value; but for Strict it is correlated. Both users, in the second statement †, suffer an “over the shoulder” attack against the new password.

These programs, although presented informally here, have a precise denotation in the model of [22]: this example is not relying on informality.

■ **Figure 1** Updating a password

1.1 Changing a password: is it only “fresh”, or actually “different”?

Consider the situation in Fig. 1 with state space a set $\{A, B, C\}$ of just three possible one-letter passwords (to keep things simple); we assume the distribution of passwords is known, and (again for simplicity) we assume it’s uniform. Now contrast two users’ behaviours when forced to change their passwords: User Lax makes a fresh, uniform choice of password; User Strict however is forced to choose a different password, still uniformly but not the one he already has. In Fig. 1 both users then suffer an “over the shoulder” attack where the adversary glimpses a letter the user did *not* type as he logged-in with his new password.

Although it’s common to consider leaks wrt. initial values, especially e.g. for channels that don’t update any state anyway, in Fig. 1 instead it’s very natural to focus on the *final* state of these two programs, i.e. the new password: the old password is not so interesting to the shoulder-cruiser. This “focus on final” obtains for the theory as well, but for a different reason: our aim (in program semantics) is to integrate security “correctness” with (ordinary) functional correctness of programs, i.e. to treat the two within the same framework [22]; and since functional correctness (and correctness comparisons, i.e. refinement [24]) is determined wrt. the final values a program produces, we should do the same for security correctness: more is said about this in §5.3 and its associated App. G.

The example above was deliberately constructed so that the two programs have the same final distribution: with our assumption of uniformly distributed passwords beforehand, in both cases (Lax, Strict) the new password’s distribution will be uniform too. Furthermore, the effect of the leak statement is also the same: the adversary has the same knowledge about the new password in both cases — he knows exactly one value that it is not. So are these programs equivalent in terms of their functional- and information-flow behaviour?

The answer depends on the context of operation. As closed systems with a single secret X , they are indeed equivalent when the initial distribution is uniform.¹ But they are not equivalent if we consider correlations between the final value of X , the new password, and some

¹ Complex examples can give equivalence for all initial distributions; but this simple case makes our point.

other variable, call it Z , even though Z is not mentioned in either program. That is because to any adversary aware of that correlation, the first program will leak information about Z but the second will not: for example, if it is known that each user has the same password for their Facebook (X) and Twitter (Z) accounts, then the over-the-shoulder attack against **Strict**'s new password leaks information about his (unchanged) Twitter password *even though his program does not access Twitter at all*.² This is essentially the Delanuis scenario presented in a programming-language context where X is the statistical database and the correlation with Z is “auxiliary information” [11] except that, unlike in the traditional presentation, ours here allows the “database” (the password) to be updated. In this programming context, we call Z a *collateral variable*, and information flow from it is *collateral leakage*.

The Dalenius phenomenon, so well known in security research, is truly remarkable as collateral leakage in rigorous reasoning about programs: what kind of algebra would invalidate an equality because of variables not occurring in either formula? In Fig. 1 are two sequential programs whose treatment of *all* variables they mention (X) is the same in terms of assignments to those variables and leaks about them; yet the two programs are not the same if collateral variables must be considered, variables the program does not even mention.

In a collateral-enabled context, the semantics of **Lax** and of **Strict** must differ.³

1.2 Our contributions

1. We extend our channel-based results from [1], based effectively on read-only programs, to show that even read/write programs' collateral leakage can be bounded without knowing what the collateral variables might be (§3),
2. and that the bounds can be used in practice, e.g. in cryptography (§4).
3. We explain the connection between collateral leakage and compositional semantics for secure programs: that the “problems” with collateral leakage, in both informal (Dalenius) and rigorous (program semantics) settings, are the same (§5.1).
4. We sketch how to extend our treatment of abstract channels [23] and abstract Hidden Markov Models [22] to construct fully abstract collateral-aware program semantics (§5.2).
5. We justify for our strong secure-refinement order, by reducing it to the simpler Bayes vulnerability [26] tests in a collateral context (§5.3).

In (§2) we provide the background for our contributions. Full proofs are in the appendix. Notation is introduced as needed, with a full glossary for reference in App. A.

2 Abstract information flow: channels, hypers, vulnerability

2.1 Abstract semantics of channels, based on hyper-distributions

Given set \mathcal{X} of inputs and \mathcal{Y} of observations, a *channel* between \mathcal{X} and \mathcal{Y} is a (stochastic) matrix whose \mathcal{X} -indexed rows sum to 1. We write the type of such channels/matrices as $\mathcal{X} \rightarrow \mathcal{Y}$ and for $C: \mathcal{X} \rightarrow \mathcal{Y}$ its constituents are elements $C_{x,y}$ at row x and column y that gives the conditional probability of output y from input x , the x 'th row $C_{x,-}$ and the y 'th column $C_{-,y}$. For any set \mathcal{S} we write $\mathbb{D}\mathcal{S}$ for the set of discrete distributions on \mathcal{S} ; thus for example

² It's interesting to reflect that many *IT* departments might impose **Strict**'s policy; but it's *his* Twitter account that is at risk, not **Lax**'s.

³ A more extreme example of this is the two programs $X := 0$ and **leak** X ; $X := 0$ where both programs simply set X to 0 and so an adversary, knowing the code, knows also that X will finally be 0. Yet the second program can reveal something about a collateral Z , while the first cannot.

any row $C_{x,-}$ of $C: \mathcal{X} \rightarrow \mathcal{Y}$ can be interpreted as an element of $\mathbb{D}\mathcal{Y}$. Similarly the $\mathcal{X} \times \mathcal{Y}$ matrix-subtype $\mathcal{X} \leftarrow \mathcal{Y}$ is co-stochastic, a channel from \mathcal{Y} to \mathcal{X} with 1-summing columns.

Following [23, 2] an *abstract* channel is a function $\mathbb{D}\mathcal{X} \rightarrow \mathbb{D}^2\mathcal{X}$ from a *prior* distribution on \mathcal{X} , i.e. of type $\mathbb{D}\mathcal{X}$, to a distribution of distributions-on- \mathcal{X} , thus of type $\mathbb{D}(\mathbb{D}\mathcal{X})$, equivalently $\mathbb{D}^2\mathcal{X}$. That latter is the type of *hyper-distributions* on \mathcal{X} , or “hypers” for short. To do that, we form the joint distribution induced by the prior and the channel, but then abstract from the observations’ values $y: \mathcal{Y}$, retaining only the effect they have on our a-posteriori, i.e. Bayesian reasoning that revises the prior. The support of the hyper is the set of posterior distributions induced on the prior by each possible observed value $y: \mathcal{Y}$, and the probability assigned by the hyper to each such posterior, i.e. each element of its support, is the (marginal) probability associated with the observation y that induced it. In more detail — given a channel $C: \mathcal{X} \rightarrow \mathcal{Y}$ and prior $\pi: \mathcal{X}$, the resulting hyper $\Delta: \mathbb{D}^2\mathcal{X}$ is found by: constructing the joint distribution $J: \mathbb{D}(\mathcal{X} \times \mathcal{Y})$ given by $J_{x,y} := \pi_x C_{x,y}$; then for each y the column $J_{-,y}$ normalising⁴ to give the posterior induced on π by that y ; and finally taking for the probability assigned to that posterior the normalising factor. We write $\pi \triangleright C$ for the joint distribution J and $[-]$ for the abstraction from y , so that $\Delta = [\pi \triangleright C]$ describes the whole hyper-construction procedure.

As an example, in Fig. 1 above the statement `leak [X+, X-]` acts, as a channel, by taking the (initial, and uniform) prior `[A, B, C]` produced by the assignment statement just before to the final hyper `[[B, C], [C, A], [A, B]]` where each of the posterior distributions is uniform (over two values), and the hyper itself is a uniform distribution on the three of them. An adversary does not know beforehand which two values afterwards he will know are possible; afterwards he does know the two possibilities, but still does not know which is right.⁵

2.2 Vulnerability induced by gain-functions

Vulnerability [2] is a generalisation of entropy (of distributions), no longer necessarily e.g. Shannon but now others more adapted for secure programming, and whose great variety allows fine-grained control of the significance of the information that might be leaked [2, 3].

Given a state-space \mathcal{X} , vulnerability is induced by a *gain function* over that space, typically g of type $\mathbb{G}_w\mathcal{X} = \mathcal{W} \rightarrow \mathcal{X} \rightarrow \mathbb{R}$, for some space of *choices* $w: \mathcal{W}$. When \mathcal{W} is obvious from context, or unimportant, we will omit it and write just $g: \mathbb{G}\mathcal{X}$. Given g and w (but not yet x) the function $g.w$ is of type $\mathcal{X} \rightarrow \mathbb{R}$ ⁶ and can thus be regarded as a random variable on \mathcal{X} . As such, it has an expected value on any distribution π over \mathcal{X} , written $\mathcal{E}_\pi g.w$.⁷

Once we have x , the (scalar) value $g.w.x$ is simply of type \mathbb{R} and represents the gain to an adversary if he guesses w when the secret’s actual value is x . This suggests that \mathcal{W} should simply be \mathcal{X} , i.e. that there would be no point in guessing a w that was not in \mathcal{X} , and indeed a particularly simple example is $\mathcal{W} = \mathcal{X}$ with $g.w.x = (1 \text{ if } w=x \text{ else } 0)$ so that the adversary gains 1 if he guesses correctly and 0 otherwise: we call this particular gain-function g_{id} . However for reasons explored thoroughly elsewhere [2, 23], there are *practical* advantages to allowing \mathcal{W} to be more general. And a *theoretical* benefit is that it is the more general \mathcal{W} ’s that allow representation of many conventional entropy functions (including even Shannon, after some encoding), thus bringing them all within the same framework [23].

⁴ If several distinct y ’s produce the same posterior, they are amalgamated; if there is y with zero marginal probability, it and its (undefined) posterior are omitted.

⁵ The first is an “unknown known” and the second is a “known unknown”.

⁶ We write dot for function application, left associative, so that function g applied to argument w is $g.w$ and then $g.w.x$ is $(g.w)$ applied to x , that is using the Currying technique of functional programming. This convention reduces clutter of parentheses, as we see later.

⁷ In general we write $\mathcal{E}_\pi f$ for the expected value of function $f: \mathcal{X} \rightarrow \mathbb{R}$ on distribution $\pi: \mathbb{D}\mathcal{X}$.

A gain function $g: \mathbb{G}\mathcal{X}$ induces a g -vulnerability function $V_g: \mathbb{D}\mathcal{X} \rightarrow \mathbb{R}$ so that $V_g[\pi]$ for $\Pi: \mathbb{D}\mathcal{X}$ is the maximum over all choices $w: \mathcal{W}$ of the expected value of $g.w$ on π , that is $\max_w (\mathcal{E}_\pi g.w)$. In the simple 1-or-0 case above, the vulnerability $V_{g_{id}}$ is called the *Bayes vulnerability*, sometimes written just V ; it is one-minus the Bayes-Risk of Decision Theory.

Vulnerability applies also to hypers, so that $V_g \Delta$ for $\Delta: \mathbb{D}^2\mathcal{X}$ is the expected value of V_g over Δ , itself considered as a distribution on $\mathbb{D}\mathcal{X}$. That allows us to write the succinct $V_g[\pi \triangleright C]$ for the a posteriori g -vulnerability of prior π through channel C .⁸

2.3 (Classical) Gain-function leakage and capacity for channels

In general, information leakage on prior π due to channel C is a comparison between the “information content” of prior knowledge about π , and of posterior knowledge after observing C ’s output; in particular, its precise definition depends on how the information is measured. Here we will use gain functions and speak of g -leakage (rather than e.g. the more specific Shannon leakage). The *multiplicative g -leakage* of C wrt prior π and gain function g is the (\log_2 of the) ratio between the posterior and prior g -vulnerabilities:

$$\mathcal{L}_g(\pi, C) := \lg(V_g[\pi \triangleright C] / V_g[\pi]) \quad . \quad (1)$$

The *capacity* of a channel is the supremum of that leakage (1), but varying in its definition depending on whether the supremum is over either gain functions, or priors or both:

$$\mathcal{L}_\forall(\pi, C) := \sup_g \mathcal{L}_g(\pi, C) \quad , \quad \mathcal{L}_g(\forall, C) := \sup_\pi \mathcal{L}_g(\pi, C) \quad , \quad \mathcal{L}_\forall(\forall, C) := \sup_{\pi, g} \mathcal{L}_g(\pi, C) \quad .$$

Remarkably, it can be shown that $\mathcal{L}_\forall(\forall, C)$ equals $\mathcal{L}_{g_{id}}(\forall, C)$ (“min-capacity”): it is the most robust estimation of leakage, and can always be achieved for some scenario and prior [2].

2.4 Collateral Gain-function leakage and capacity for channels

The \mathcal{Z} -collateral leakage of a channel $C: \mathcal{X} \rightarrow \mathcal{Y}$ is with respect to some third space \mathcal{Z} and is induced by knowledge of a correlation between values in \mathcal{X} and \mathcal{Z} .⁹ We write $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$ of resp. types $\mathcal{X}, \mathcal{Y}, \mathcal{Z}$ equivalently for random- or program variables, depending on context. We assume a known correlation between \mathbf{X} and a fresh secret \mathbf{Z} given by a joint probability distribution $\Pi: \mathbb{D}(\mathcal{Z} \times \mathcal{X})$ so that leaks about \mathbf{X} through (a computer program implementing) channel C induce “collateral” leaks about \mathbf{Z} as follows. Define $\tilde{\pi}$ to be the \mathcal{Z} -marginal of Π and determine some channel $\tilde{\Pi}$ in $\mathcal{Z} \rightarrow \mathcal{X}$ so that $\Pi = \tilde{\pi} \triangleright \tilde{\Pi}$: this factors Π into its marginal and a conditional.¹⁰ Then the matrix multiplication $\tilde{\Pi} \cdot C$ gives a channel D (for Delanius) of type $\mathcal{Z} \rightarrow \mathcal{Y}$, and the collateral leakage of C is the ordinary leakage of D , calculated as at (1) above. (The (matrix) multiplication $\tilde{\Pi} \cdot C$ is the *channel cascade* of $\tilde{\Pi}$ and C .)

The \mathcal{Z}/Π -collateral capacity of channel C is then the upper bound of the above with respect to all gain functions, that is $\mathcal{L}_\forall(\tilde{\pi}, \tilde{\Pi} \cdot C)$. The collateral capacity of channel C in general is then taken over all possible \mathcal{Z} and Π .

⁸ Recalling that $[-]$ around a list of one element makes the point distribution on that element, we can regard $V_g[\pi]$ as the expected value of V_g on the point hyper on π , agreeing with our usage just above.

⁹ In Dalenius terms \mathcal{X} is the statistical database, with queries of type \mathcal{Y} , and the correlation between \mathcal{X} and \mathcal{Z} is auxiliary information [11].

¹⁰ Recall that $\Pi_{z,x}$ is thus $\tilde{\pi}_z \tilde{\Pi}_{z,x}$. Unless marginal $\tilde{\pi}$ is full support, conditional $\tilde{\Pi}$ is not unique. But it does not matter: any choice suffices, and does not affect the outcome.

3 Collateral gain-function leakage and capacity for *HMM*'s

3.1 *HMM*'s as models for leaking probabilistic programs

The discussion in §2, for channels only, corresponds to considering read-only programs. With *Hidden Markov Models* we combine channels and Markov transitions between initial- and final states (e.g. over variables \mathbf{X} in \mathcal{X}), thus modelling read/write programs; and the constructions of §2 carry over. A Markov transition is described by a matrix M say so that $M_{x,x'}$ is the probability that initial state x will result in final state x' . An *HMM-step* then comprises a channel and a transition together, but acting independently on the initial state: we call C its channel and M *markov* (lower case), and write it $(C:M)$ of type $\mathcal{X} \rightarrow \mathcal{Y} \times \mathcal{X}$. Defined $(C:M)_{x,y,x'} = C_{x,y} M_{x,x'}$, it is a single stochastic matrix with rows \mathcal{X} and columns $\mathcal{Y} \times \mathcal{X}$.

HMM-steps have the special characteristic that their C and M effects (i.e. output values y, x') are independent (uncorrelated) for each separate input-value x ; but two steps, say $H^{1,2}$ of the same type $\mathcal{X} \rightarrow \mathcal{Y} \times \mathcal{X}$ can be *sequentially composed* to give a single *HMM* again: this is natural if we are using them to model programs. The composed type is the $\mathcal{X} \rightarrow \mathcal{Y}^2 \times \mathcal{X}$ that takes initial state x to final state x' via some intermediate state x'' , leaking information (y^1, y^2) – as it goes – gradually into the set \mathcal{Y}^2 . We define

$$(H^1; H^2)_{x,(y^1,y^2),x'} = \sum_{x''} H_{x,y^1,x''}^1 H_{x'',y^2,x'}^2, \quad (2)$$

and note that it is again stochastic, but no longer necessarily a *step* — its outputs y^2 and x' might be correlated even for a single input x . Because of that, sequential compositions are strictly more general than the *HMM*-steps built directly from $(C:M)$ — that is, for $C^{1,2}, M^{1,2}$ in general there is not necessarily a single C, M such that $(C^1:M^1); (C^2:M^2) = (C:M)$. In App. E we give some elementary properties of *HMM*-steps, showing in particular how the different encodings of markovs and channels preserve their different purposes and allows the same definition of sequential composition to be used for both.

More generally, two *HMM*'s over the same state \mathcal{X} but distinct observation spaces $\mathcal{Y}^{1,2}$ can be composed as just above; in that case the composite observation space is $\mathcal{Y}^1 \times \mathcal{Y}^2$, and that is what allows us to use them for modelling programs which, in general, are not simply just a fixed sequence of assignment statements. A monadic account of this generality, including loops and conditionals, is given in [21].

3.2 The induced channel of an *HMM*

Recall that an *HMM* of type $\mathcal{X} \rightarrow \mathcal{Y} \times \mathcal{X}$ takes an initial state in \mathcal{X} to a final state (also) in \mathcal{X} and leaks observations in \mathcal{Y} along the way, where “along the way” is more general than channels: in Fig. 1 for example the `leak` statement occurred *after* a Markov update. Collateral leakage however, as we saw in the previous section, is with respect to a value Z which the *HMM* (or program) does not even mention: it neither reads nor updates it; the programmer has never even heard of it. But it is somehow correlated with the initial state: that being so, we show how to isolate the channel part of an *HMM*, combining all its leaks and ignoring its final state. Fix H as some *HMM* of type $\mathcal{X} \rightarrow \mathcal{Y} \times \mathcal{X}$ below.

► **Definition 1.** The *effective channel* of H , written $\text{chn}.H$, is a stochastic matrix of type $\mathcal{X} \rightarrow \mathcal{Y}$ and defined simply by ignoring the final state: thus $(\text{chn}.H)_{x,y} := \sum_{x'} H_{xyx'}$.

► **Definition 2.** The collateral leakage resp. capacity of H wrt a prior $\Pi: \mathbb{D}(\mathcal{Z} \times \mathcal{X})$ is the collateral leakage resp. capacity of $\vec{\Pi} \cdot \text{chn}.H$. (This is well-defined, though $\vec{\Pi}$ is not unique.)

The simplicity of Defns. 1,2 conceals that it can be difficult in practice to calculate the collateral leakage of an *HMM*, since Def. 1 mandates calculating the whole H first, only then abstracting from its final state. Yet if H is expressed as a sequential composition of many smaller ones, e.g. we have $H = H^1; H^2; \dots; H^N$, still the final states of the *intermediate* H^n 's must be retained, not only to form the composition, but because the overall y observation from H comprises all the smaller observations $y^1 \dots y^N$ with each y^{n+1} being determined by the final state $(x')^n$ of the H^n just before — we can abstract only at the very end.

In the special case however where each H^n is an *HMM*-step $(C^n:M^n)$, the calculation of the effective channel can be somewhat decomposed.

► **Lemma 3.** *Let H be an *HMM* and $(C:M)$ an *HMM*-step. Then* [proof in App. C]

$$\begin{aligned} \text{chn.}(C:M) &= C \\ \text{chn.}((C:M); H) &= C \parallel (M \cdot \text{chn.}H) \end{aligned}$$

where in general $(C^1 \parallel C^2)_{x, (y^1, y^2)} = C^1_{x, y^1} C^2_{x, y^2}$ is parallel composition of channels. The M cannot be discarded, since it affects the prior of the “tail” H of the sequential composition.

Even with Lem. 3, in general $\text{chn.}H$ can be challenging to compute because its size (given by the number of columns in the stochastic matrix representation) grows exponentially with the number of single-step *HMM*'s, in the definition of H , that have non-trivial channel portions. We give an example of such a calculation in §4 (fast exponentiation for cryptography).

On the other hand, if we want to compute only the collateral *capacity*, we can obtain at least an upper bound at considerably less cost, without the need to compute $\text{chn.}H$ exactly. The following provides an upper bound for $\mathcal{L}_\forall(\forall, \text{chn.}H)$, and requires only linear resources.

► **Lemma 4.** *For any H let $\text{CCap.}H$ be defined* [proof in App. C]

$$\begin{aligned} \text{CCap.}(C:M) &= \mathcal{L}_\forall(\forall, C) && \text{if } H=(C:M) \\ \text{CCap.}((C:M); H') &= \mathcal{L}_\forall(\forall, C) + \min(\mathcal{L}_\forall(\forall, M), \text{CCap.}H') && \text{if } H=(C:M); H' \end{aligned}$$

Then $\mathcal{L}_\forall(\forall, \text{chn.}H) \leq \text{CCap.}H$ with the stochastic matrix M (on rhs) treated as a channel.

In fact Lem. 4 provides a very robust estimate of the collateral capacity of an *HMM*, since it does not mention \mathcal{Z} or the correlating Π . And it is the best possible general bound, achieving equality for some examples, e.g. Fig. 1. It is also easy to calculate since for any channel we have from [2] that $\mathcal{L}_\forall(\forall, C) = \mathcal{L}_{g_{id}}(\Upsilon_{\mathcal{X}}, C)$, where $\Upsilon_{\mathcal{X}}$ is the uniform prior on \mathcal{X} .

On the other hand, the bound can be very conservative; we provide an example in App. C.

In cases where we know the correlation Π (and thus \mathcal{Z}), we can compute the collateral capacity by identifying the gain function that achieves the extremal value in Def. 2.

► **Theorem 5.** *Given H and $\Pi: \mathbb{D}(\mathcal{Z} \times \mathcal{X})$ with $\tilde{\pi}, \vec{\pi}$ resp. the marginals of Π on \mathcal{Z}, \mathcal{X} ; define conditional $\vec{\Pi}$ as in §2.4. There exists $\hat{g}: \mathbb{G}_{\mathcal{Z}} \mathcal{Z}$ and $\hat{g}^\Pi: \mathbb{G}_{\mathcal{Z}} \mathcal{X}$ such that* [proof in App. C]

$$\mathcal{L}_\forall(\tilde{\pi}, \vec{\Pi} \cdot \text{chn.}H) = \mathcal{L}_{\hat{g}}(\tilde{\pi}, \vec{\Pi} \cdot \text{chn.}H) = \mathcal{L}_{\hat{g}^\Pi}(\tilde{\pi}, \text{chn.}H) .$$

This shows that it is possible to construct the gain-function that maximizes the collateral capacity, and allows its exact calculation. Moreover, it also shows that the collateral capacity of H wrt. \mathcal{Z} can be understood as regular g -leakage of H wrt. the *initial* state of \mathcal{X} .

The next theorem is more general, and gives an upper bound over all possible correlations: it is determined by the extremal leakage of the initial prior $\pi: \mathbb{D}\mathcal{X}$, thus easy to calculate [1].

► **Theorem 6.** *Given H and Π as above, $\Upsilon_{\mathcal{X}}$ is uniform on \mathcal{X} , then* [proof in App. C]

$$\mathcal{L}_\forall(\tilde{\pi}, \vec{\Pi} \cdot \text{chn.}H) \leq \mathcal{L}_{g_{id}}(\Upsilon_{\mathcal{X}}, \text{chn.}H) , \quad \text{where } \tilde{\pi}, \vec{\Pi} \text{ are as defined in Thm. 5.}$$

```

// B for base, the cleartext; E for exponent, the key: precondition is B,E >= 0,0 .
// P for power, the ciphertext.
P:= 1
while E!=0 // Invariant is  $P \cdot (B^E) = b^e$ , where  $b,e$  are initial values of B,E .
  D:∈ [2,3,5] // D for divisor; uniform choice from {2,3,5}.
  R:= E mod D; // R for remainder.
  if R!=0 then P:= P*B^R fi // Side-channel: is E divisible exactly by D ?
  B:= B^D // D is small: assume no side-channel here.
  E:= E div D // State update of E here. (No side-channel.)
end
// Now  $P=b^e$  and E=0: but what has an adversary learned about the initial  $e$  ?

```

Although our state comprises B,E,P,D,R we concentrate only on the secrecy of E. In particular, we are not trying to discover B or P in this case; and D,R are of no external significance afterwards anyway.

■ **Figure 2** Defence against side channel analysis in exponentiation

Note that when X,Z are completely correlated, i.e. when $\vec{\Pi}$ and $\vec{\Pi}$ are both the identity, the inequality in Thm. 6 becomes equality.

4 Example: Collateral damage from leaky cryptography

Keys for public-key cryptography are best if independent; but recently [17] discovered an unexpected sharing of the prime numbers used to generate them. Thus information leaked while using one secret key to encrypt a message, even if the message itself remains secure, could cause collateral leakage wrt. some other key elsewhere and so put future encryptions at risk — even if they are at another site apparently having no connection to the first. That motivates our example here, the collateral leakage from a fast-exponentiation algorithm.

Fig. 2 implements fast exponentiation with a random choice of divisor that defends against a side channel that leaks program flow (of a conditional) [28]. Since the program code is public, that leak is effectively of whether divisor D exactly divides the current value of E, which value is steadily decreased by the update at the end of each iteration: thus additional information is leaked every time. In the standard (and fastest) algorithm the divisor D is always 2, but that ultimately leaks E’s initial value exactly, one bit (literally) on each iteration. The final value of E is always zero, of no significance; but its initial value represents collateral leakage about subsequent use of this same key (obviously), but also the use of other apparently unrelated keys elsewhere [17]. The obfuscating defence is to vary the choice of D unpredictably from one iteration to the next, choosing it secretly from some set \mathcal{D} , here $\{2, 3, 5\}$ although other options are possible. The divisor D itself is not leaked.

We modelled the loop as a sequential composition of *HMM*-steps for a fixed number of iterations and used Lem. 3 to construct a channel that captures the leakage of information about the initial state of the program. Although our calculation is wrt. the uniform prior on E, and even though we do not know the extent of any correlation between this key E and others used elsewhere, by using the “Miracle Theorem” [2, Thm 5.1], we can bound the maximum leakage about the initial value of E with the min-capacity of such a channel. Furthermore, by relying on Thm. 6 we can see that this min-capacity can also be used as a bound on the collateral leakage *with respect to any other secret that might be correlated to E*.

Detailed calculations are in App. D. Our Tbl. 1 confirms that the larger the divisor set \mathcal{D} , the less effective is the side channel; and the protection is increased with more bits for E.

Size of E	$\mathcal{D}=\{2\}$	$\mathcal{D}=\{2, 3\}$	$\mathcal{D}=\{2, 3, 5\}$
4 bits	4	2.80	2.22
5 bits	5	3.32	2.61
6 bits	6	3.83	2.92
7 bits	7	4.34	3.21
8 bits	8	4.88	3.51

Note that in the case $\mathcal{D}=\{2\}$ the whole secret E is leaked.

As explained at end §2.3, that \mathcal{L}_{gid} gives the upper bound \mathcal{L}_\forall for all vulnerabilities.

■ **Table 1** Collateral leakage $\mathcal{L}_{gid}(\forall, Prog)$ in bits wrt E for different \mathcal{D} 's, for $Prog$ given at Fig. 2.

5 The abstract semantics of “collateral aware” *HMM*'s

5.1 Bits-leaked vs. a partial order on information-flow

We now contrast our practical example in §4, just above, with our earlier treatment of passwords, in particular our description (at the end of §2.1) of leakage abstractly as a hyper-distribution rather than as a number of bits (Tbl. 1).

Both password programs take prior $[A, B, C]$ to the hyper $[[B, C], [C, A], [A, B]]$ on the final state, as we saw in §2.1. But for collateral leakage we must consider (also) the initial state: and here the two programs differ. The hyper representing our knowledge of the prior distribution of the *initial* password, after execution of *Lax*, is still $[[A, B, C]]$ — it expresses that we know for sure (point hyper $[[\cdot \cdot \cdot]]$) that the distribution of X is uniform ($[A, B, C]$), that we have learned nothing (more) about the distribution of the password initially. But for program *Strict* the hyper describing our revised knowledge about the initial password is

$$[[A^{\oplus 1/2}, B^{\oplus 1/4}, C^{\oplus 1/4}], [B^{\oplus 1/2}, C^{\oplus 1/4}, A^{\oplus 1/4}], [C^{\oplus 1/2}, A^{\oplus 1/4}, B^{\oplus 1/4}]] , \quad (3)$$

now a uniform hyper over three, different, *skewed* inners: the first $[A^{\oplus 1/2}, B^{\oplus 1/4}, C^{\oplus 1/4}]$, for example, is the distribution A wprob (with probability) $1/2$ and B, C wprob $1/4$ each.¹¹ An adversary, knowing the program text of *Strict* and having calculated this hyper would conclude that, with probability $1/3$ each, he can after the run of the program revise his knowledge of the initial-password distributions from uniform to one of these three skewed inners (although he does not know beforehand which one of those three possible revisions will eventuate).

Abstracting from any particular vulnerability (or entropy), we conclude that although *Lax* and *Strict* are equivalent wrt their setting and security of the new password when the previous password was uniformly distributed (§1.1), wrt collateral security *Strict* is less secure than *Lax*. This abstract conclusion, i.e. reached without calculating specifics of “number of bits leaked”, comes from the “no less secure than” partial order (\sqsubseteq) on hypers $\mathbb{D}^2\mathcal{X}$ where indeed (3) $\sqsubset [[A, B, C]]$, that is the *lhs*-hyper is strictly less secure than the *rhs*-hyper [19, 2, 21, 1, 23, 22]. The significance is that the theory associated with (\sqsubseteq) shows that for *any* hypers $\Delta^{1,2}$ we have $\Delta^1 \sqsubseteq \Delta^2$ exactly when $V_g(\Delta^1) \geq V_g(\Delta^2)$ for *all* gain-functions g : increased security corresponds to decreased vulnerability [1]. Since gain functions can express not only Bayes Vulnerability but also Shannon entropy (with some encoding) and many other entropies besides, that is a very strong result: for any “favourite entropy” you might use to compare the security of two programs, if it is expressible as V_g for some g (no matter how bizarre) then you will find the former at least as secure as the latter.

¹¹ The details of calculating this hyper are given in App. B. Intuitively, whichever (final) not-password is seen by the adversary, it is more than uniformly probable that the initial password was that since, by *Strict*'s policy, the old password cannot be the new password that the adversary did not see.

For example, the Shannon entropy of $[[A, B, C]]$ is $\lg 3 \approx 1.58$; of (3) it is $1/2 \cdot 1 + 1/4 \cdot 2 + 1/2 \cdot 2 = 1.5$. (Decreasing Shannon entropy indicates increasing vulnerability.) And the Bayes Vulnerability of $[[A, B, C]]$ is $1/3$; of (3) it is $2/3$, increasing as expected.

5.2 Collateral-aware fully-abstract semantics of *HMM*'s

All proofs for this section are to be found in App. F.

We saw in §2.1 an abstract semantics for (read only) *channels* in $\mathbb{D}\mathcal{X} \rightarrow \mathbb{D}^2\mathcal{X}$; we could summarise that by writing $[[C]].\pi = [\pi \triangleright C]$, that is that $[[\cdot]]$ is the semantic function that takes a channel matrix in $\mathcal{X} \rightarrow \mathcal{Y} \times \mathcal{X}$ to a function as just above. For *HMM*'s we do similar:

► **Definition 7.** Given $H: \mathcal{X} \rightarrow \mathcal{Y} \times \mathcal{X}$ its denotation in $\mathbb{D}\mathcal{X} \rightarrow \mathbb{D}^2(\mathcal{X} \times \mathcal{X}')$ takes prior π to the hyper $[\pi \triangleright H]$ where for hypers (generalising channels) we define $(\pi \triangleright H)_{x,y,x'} = \pi_x H_{x,y,x'}$, and using \mathcal{Y} -abstraction now (effectively) on $\mathcal{Y} \times (\mathcal{X} \times \mathcal{X}')$ we obtain a hyper in $\mathbb{D}^2(\mathcal{X} \times \mathcal{X}')$. (The state-spaces \mathcal{X} and \mathcal{X}' are the same: we use the prime only to distinguish final and initial.) We write $[[H]]$ for the resulting function, and call those $[[\cdot]]$ -images *abstract HMM's*.

Def. 7 extends our earlier definition of *HMM*-semantics [22] by retaining the initial state, so that we can account for collateral effects.¹² This new type is the bottom arrow of Fig. 3.

For any abstract *HMM* h and any collateral type \mathcal{Z} we can then define a “ \mathcal{Z} -lifted” function $h^{\mathcal{Z}}: \mathbb{D}(\mathcal{Z} \times \mathcal{X}) \rightarrow \mathbb{D}^2(\mathcal{Z} \times \mathcal{X}')$ that takes a collateral correlation $\Pi: \mathbb{D}(\mathcal{Z} \times \mathcal{X})$ to a hyper in $\mathbb{D}^2(\mathcal{Z} \times \mathcal{X}')$ based on such correlations: this is the top arrow of Fig. 3. We then have

► **Definition 8.** For abstract *HMM*'s $h_{1,2}$ define $h_1; h_2$ to be $h_1; h_2^{\mathcal{X}}$ where (bold) $;$ is forward Kleisli composition on the (larger) state space $\mathcal{X} \times \mathcal{X}'$, as used for monadic *HMM*'s [22], and $-\mathcal{X}$ is as in Fig. 3 but taking \mathcal{Z} to be \mathcal{X} . (Recall that \mathcal{X} and \mathcal{X}' are the same.) It satisfies the essential property that $[[H^1; H^2]] = [[H^1]]; [[H^2]]$ (Lem. 13).

Vulnerabilities V_g are similarly extended, now taking $g: \mathcal{W} \rightarrow (\mathcal{X} \times \mathcal{X}') \rightarrow \mathbb{R}$, i.e. referring to both initial and final states. With them, we define *secure refinement* for *HMM*'s, that $h_1 \sqsubseteq h_2$ just when $V_g(h_1.\pi) \geq V_g(h_2.\pi)$ for all such g 's and $\pi: \mathbb{D}\mathcal{X}$; and *behavioural equivalence* to be refinement in both directions. We say that (concrete) *HMM*'s are behaviourally equivalent just when their (abstract) denotations are. Because refinement \sqsubseteq is anti-symmetric [19, 1] that establishes full abstraction of $[[\cdot]]$ for that equivalence on concrete *HMM*'s. A basic compositionality result is then that Def. 8 respects this equivalence (Thm. 14).


A second basic compositionality result concerns state extension: an *HMM* acting on state \mathcal{X} , i.e. on a single variable \mathbf{X} , can be also considered to be acting on part of a larger state $\mathcal{Z} \times \mathcal{X}$ with an extra \mathbf{Z} (whether correlated or not). Even if the program does not read or write \mathbf{Z} , it nevertheless can release information about \mathbf{Z} : this is the essence of the collateral “problem”. The two programs **leak** \mathbf{X} ; $\mathbf{X} := 0$ and simply $\mathbf{X} := 0$ make it clear that the final-state semantics is insufficient for the extension: those two programs both take any $\pi: \mathbb{D}\mathcal{X}$ to the same final hyper $[[0]]$, that \mathbf{X} is “certainly 0”. Yet in the context $\mathbf{X} := \mathbf{Z}; (-)$ they are different: one reveals \mathbf{Z} and the other does not.

► **Definition 9.** Given abstract h we define its \mathcal{Z} extension $h^{\times \mathcal{Z}}: \mathbb{D}(\mathcal{Z} \times \mathcal{X}) \rightarrow \mathbb{D}^2(\mathcal{Z} \times \mathcal{X})^2$ in two steps.¹³ First, we construct $h^{\mathcal{X} \times \mathcal{Z}^2}$ as in Fig. 3 to produce function a function of type $\mathbb{D}(\mathcal{Z}^2 \times \mathcal{X}^2) \rightarrow \mathbb{D}^2(\mathcal{Z} \times \mathcal{X})^2$ that carries the initial value of \mathcal{X} as well as two copies of \mathcal{Z} .

¹² Our earlier domain was $\mathbb{D}\mathcal{X} \rightarrow \mathbb{D}^2\mathcal{X}$, as is still the case for channels since do not update their state.

¹³ We now drop the \mathcal{X}' primes, to reduce clutter.

At left: Given a collateral correlation $\Pi: \mathbb{D}(\mathcal{Z} \times \mathcal{X})$ let the induced collateral channel $Z: \mathcal{Z} \leftarrow \mathcal{X}$ and \mathcal{X} -prior $\pi: \mathbb{D}\mathcal{X}$ be such that $\Pi = Z(\pi)$, that is so that $\Pi_{z,x} = Z_{z,x}\pi_x$.



At right (in grey): For an initial/final distribution $\delta: \mathbb{D}(\mathcal{X} \times \mathcal{X}')$, the left-multiplication $Z \cdot \delta$ produces a distribution in $\mathbb{D}(\mathcal{Z} \times \mathcal{X}')$, just as matrix multiplication would.

At right (in black): The \mathbb{D} -lifting (push forward) of the multiplication $(Z \cdot)$ thus takes an initial-final hyper in $\mathbb{D}^2(\mathcal{X} \times \mathcal{X}')$ to a hyper in $\mathbb{D}^2(\mathcal{Z} \times \mathcal{X}')$.

Summary: A collateral Z is linked to our state X by joint distribution $\Pi: \mathbb{D}(\mathcal{Z} \times \mathcal{X})$. This Π can be decomposed into its right marginal $\pi: \mathcal{X}$ on our state space, and a collateral channel $Z: \mathcal{Z} \leftarrow \mathcal{X}$ between it and Z , i.e a right conditional of Π . (Right-conditionals are not necessarily unique; but the variation on x 's where $\pi.x=0$ does not affect $\mathbb{D}(Z \cdot)$ at right (Lem. 11).

Although our abstract h at bottom can “see” only X and prior π , it can have a collateral effect on Z , given by the “ Z -lifting” h^Z shown at top. Note that h^Z is not itself an abstract HMM.

■ **Figure 3** Collateral correlation $\Pi: \mathbb{D}(\mathcal{Z} \times \mathcal{X})$ lifts $\mathbb{D}\mathcal{X} \rightarrow \mathbb{D}^2(\mathcal{X} \times \mathcal{X}')$ to $\mathbb{D}(\mathcal{Z} \times \mathcal{X}) \rightarrow \mathbb{D}^2(\mathcal{Z} \times \mathcal{X}')$.

Note that we re-order the set arguments of the \mathbb{D} 's up to commutativity. Then we define the “duplicator” $\zeta: \mathcal{Z} \times \mathcal{X} \rightarrow (\mathcal{Z} \times \mathcal{X})^2$ as $\zeta.(z, x) = (z, x, z, x)$, and conclude with $h^{\times Z}$ being $h^{Z^2 \times \mathcal{X}} \circ \mathbb{D}\zeta$, which is of type $\mathbb{D}(\mathcal{Z} \times \mathcal{X}) \rightarrow \mathbb{D}^2(\mathcal{Z} \times \mathcal{X})^2$ as required.

Our principal result (proved in App. F) is then

► **Theorem 10.** Let $h_{1,2}: \mathbb{D}\mathcal{X} \rightarrow \mathbb{D}^2\mathcal{X}^2$. Then $h_1 \sqsubseteq h_2$ iff $h_1^{\times Z} \sqsubseteq h_2^{\times Z}$ for every extension Z .

It is the *only if* of this theorem that shows the extended semantics $\mathbb{D}\mathcal{X} \rightarrow \mathbb{D}^2\mathcal{X}^2$ is enough to allow Z -extension, i.e. that our semantics is compositional even in collateral contexts.

5.3 Justifying REFINEMENT DENIED due to collateral g -vulnerability

The definition in §5.1 of secure-program refinement \sqsubseteq looks very strong, mandating a decrease of g -vulnerability for *all* gain-functions g . Why is it not *too* strong?

A key result of [19] was that \sqsubseteq is indeed not too strong in closed systems, provided one accepts Bayes vulnerability as a reasonable security measure [26], and compositionality as reasonable property of program semantics.¹⁴ The argument was that if some *general* vulnerability V_g , no matter how bizarre, mandates $P \not\sqsubseteq Q$ for two programs P, Q (i.e. **REFINEMENT DENIED**), then there will be a context \mathcal{C} and a prior π such that for that π the *Bayes* vulnerability of $\mathcal{C}(P)$, which is not bizarre at all, is strictly less than that of $\mathcal{C}(Q)$. In [19] the distinguishing g was used to construct a Markov update M^g acting on P, Q 's final state, based on that g , so that $P; M^g$ and $Q; M^g$ were distinguished by the simpler Bayes vulnerability V_{gid} .

But our programs **Lax,Strict** from Fig. 1 are not distinguished by any g in a closed system: only in an open system, with possible collateral variable Z , do we have **Lax** $\not\sqsubseteq$ **Strict**. That means the technique of [19] will not work here, because no g acting on \mathcal{X}' , the final state

¹⁴ Compositionality is often expressed $\llbracket P \rrbracket = \llbracket Q \rrbracket \Rightarrow \llbracket \mathcal{C}(P) \rrbracket = \llbracket \mathcal{C}(Q) \rrbracket$; but when an (antisymmetric) refinement order is available we refer to the stronger criterion of monotonicity of $\llbracket \mathcal{C}(-) \rrbracket$ wrt. that order.

alone, can distinguish the two programs: they produce the same hyper on \mathcal{X}' . (Recall \dagger in §5.1.) To make an analogous argument we instead use a collateral-aware g on $\mathcal{X} \times \mathcal{X}'$ to construct a collateral correlation Π^g between $\mathcal{Z}=\mathcal{W}$ and \mathcal{X} ; it is then the presence of the induced collateral variable W that allows us to deny the refinement. We can then argue that indeed **Lax** should not be refined by **Strict**, because in the presence of correlation Π^g between W and Z that refinement can be denied on the basis of Bayes vulnerability acting on \mathcal{W} alone.

We give more details in App. G.

6 Related work; conclusions

Classical analyses of quantitative information flow assume that the secret does not change, and early approaches to measuring insecurities in programs are based on determining a “change in uncertainty” of some “prior” value of the secret — although how to measure the uncertainty differs in approach. For example Clark et al [6] use Shannon entropy to estimate the number of bits being leaked; and Clarkson et al [8] model a change in belief. The role of capacity when the prior is not known was stressed by Chatzikokolakis et al [4]. Smith [26] demonstrated the importance of using measures which have some operational significance, and this idea was developed further [2] by introducing the notion of g -leakage to express leakage wrt. very general contexts. The partial order used here on HMM ’s is the same as the g -leakage order introduced by Alvim et al [2], but it also appeared in even earlier work [19]. Its properties have been studied extensively [1].

More recently Marzdiel et al [18] have analysed information flow of dynamic secrets using a model based on probabilistic automata. This reflects a view that in computing systems secrets are not necessarily static; our work [19, 22] addresses this idea by providing semantic domains for programs where secrets can change.

Clark et al [7] give techniques for static analysis of quantitative information flow based on Shannon entropy for a small while-language. Extended HMM ’s for modelling side channels have been explored by Karlof and Wagner [16] and Green et al [14] for e.g. key recovery. In §4 our quantitative capacity bounds on side channels are valid even for collateral leakage.

The abstract treatment of probabilistic systems with the introduction of a “refinement order” was originally due to Jones and Plotkin [15]; the ideas were extended to include demonic nondeterminism (as well as probability) by us [25]. In both cases the order (on programs) corresponds to an order determined by averaging over “probabilistic predicates” which are random variables over the state space. The first compositional refinement order for information flow appeared in [19] for security programs expressed in a simple programming language and in [1] for a channel model of information flow.

Prospects and future work. In [22] we noted that collateral leakage has a significant impact on abstract-program semantics for quantitative information flow; in this paper we have defined a novel collateral-aware HMM -style model (Def. 8) that properly takes it into account. The significance of the abstract models in this paper is that even though information flow becomes apparent *only* in a wider context of collateral leakage, our abstract semantics Def. 8 and Thm. 10 together show that programs *can still* be compared using only the variables they actually declare in their source code.

An analogue of Thm. 10 for qualitative information flow appeared in Chen’s thesis [5] where *compositional refinement* checking for two programs is reduced to comparisons involving only the declared variables. The benefit of this theorem can be seen in Chen’s automation of refinement checking for qualitative information flow. Our general result Thm. 10 holds out the prospect of similar automation for quantitative-security refinement checking.

References

- 1 Mário S. Alvim, Konstantinos Chatzikokolakis, Annabelle McIver, Carroll Morgan, Catuscia Palamidessi, and Geoffrey Smith. Additive and multiplicative notions of leakage, and their capacities. In *IEEE 27th Computer Security Foundations Symposium, CSF 2014, Vienna, Austria, 19-22 July, 2014*, pages 308–322. IEEE, 2014.
- 2 Mário S. Alvim, Kostas Chatzikokolakis, Catuscia Palamidessi, and Geoffrey Smith. Measuring information leakage using generalized gain functions. In *Proc. 25th IEEE Computer Security Foundations Symposium (CSF 2012)*, pages 265–279, June 2012.
- 3 Mário S. Alvim, Andre Scedrov, and Fred B. Schneider. When not all bits are equal: Worth-based information flow. In *Proc. 3rd Conference on Principles of Security and Trust (POST 2014)*, pages 120–139, 2014.
- 4 Konstantinos Chatzikokolakis, Catuscia Palamidessi, and Prakash Panangaden. Anonymity protocols as noisy channels. *Inf. Comput.*, 206(2-4):378–401, 2008.
- 5 Chris Chen. Evincing practical and theoretical consequences of the Shadow Semantics, October 2012. Undergraduate thesis, University of New South Wales. www.cse.unsw.edu.au/~carrrollm/CHEN_9945_shadow_semantics.pdf.
- 6 David Clark, Sebastian Hunt, and Pasquale Malacaria. Quantitative analysis of the leakage of confidential data. *Electr. Notes Theor. Comput. Sci.*, 59(3):238–251, 2001.
- 7 David Clark, Sebastian Hunt, and Pasquale Malacaria. Quantified interference for a while language. *Electr. Notes Theor. Comput. Sci.*, 112:149–166, 2005.
- 8 Michael R. Clarkson, Andrew C. Myers, and Fred B. Schneider. Belief in information flow. In *18th IEEE Computer Security Foundations Workshop, (CSFW-18 2005), 20-22 June 2005, Aix-en-Provence, France*, pages 31–45, 2005.
- 9 T. Dalenius. Towards a methodology for statistical disclosure control. *Statistik Tidskrift*, 15:429–44, 1977.
- 10 E.W. Dijkstra. *A Discipline of Programming*. Prentice-Hall, 1976.
- 11 Cynthia Dwork. Differential privacy. In *Proc. 33rd International Colloquium on Automata, Languages, and Programming (ICALP 2006)*, pages 1–12, 2006.
- 12 Barbara Espinoza and Geoffrey Smith. Min-entropy as a resource. *Information and Computation (Special Issue on Information Security as a Resource)*, 226:57–75, April 2013.
- 13 M. Giry. A categorical approach to probability theory. In *Categorical Aspects of Topology and Analysis*, volume 915 of *Lecture Notes in Mathematics*, pages 68–85. Springer, 1981.
- 14 P. J. Green, Richard Noad, and Nigel P. Smart. Further Hidden Markov model cryptanalysis. In *Cryptographic Hardware and Embedded Systems - CHES 2005, 7th International Workshop, Edinburgh, UK, August 29 - September 1, 2005, Proceedings*, pages 61–74, 2005.
- 15 C. Jones and G. Plotkin. A probabilistic powerdomain of evaluations. In *Proceedings of the IEEE 4th Annual Symposium on Logic in Computer Science*, pages 186–95, Los Alamitos, Calif., 1989. Computer Society Press.
- 16 Chris Karlof and David Wagner. Hidden Markov Model cryptanalysis. In *Cryptographic Hardware and Embedded Systems - CHES 2003, 5th International Workshop, Cologne, Germany, September 8-10, 2003, Proceedings*, pages 17–34, 2003.
- 17 Arjen K. Lenstra, James P. Hughes, Maxime Augier, Thorsten Kleinjung, and Christophe Wachter. Ron was wrong, Whit is right. Technical report, EPFL IC LACAL, Station 14, CH-1015 Lausanne, Switzerland, 2012.
- 18 Piotr Mardziel, Mário S. Alvim, Michael W. Hicks, and Michael R. Clarkson. Quantifying information flow for dynamic secrets. In *2014 IEEE Symposium on Security and Privacy, SP 2014, Berkeley, CA, USA, May 18-21, 2014*, pages 540–555, 2014.
- 19 Annabelle McIver, Larissa Meinicke, and Carroll Morgan. Compositional closure for Bayes Risk in probabilistic noninterference. In *Automata, Languages and Programming, 37th*

- International Colloquium, ICALP 2010, Bordeaux, France, July 6-10, 2010, Proceedings, Part II*, pages 223–235, 2010.
- 20 Annabelle McIver, Larissa Meinicke, and Carroll Morgan. A Kantorovich-monadic power-domain for information hiding, with probability and nondeterminism. In *Proc. LiCS 2012*, 2012.
 - 21 Annabelle McIver, Larissa Meinicke, and Carroll Morgan. Hidden-Markov program algebra with iteration. *Mathematical Structures in Computer Science*, 2014.
 - 22 Annabelle McIver, Carroll Morgan, and Tahiry Rabehaja. Abstract Hidden Markov Models: a monadic account of quantitative information flow. In *Proc. LiCS 2015*, 2015.
 - 23 Annabelle McIver, Carroll Morgan, Geoffrey Smith, Barbara Espinoza, and Larissa Meinicke. Abstract channels and their robust information-leakage ordering. In Martín Abadi and Steve Kremer, editors, *Principles of Security and Trust - Third International Conference, POST 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014, Proceedings*, volume 8414 of *Lecture Notes in Computer Science*, pages 83–102. Springer, 2014.
 - 24 C.C. Morgan. *Programming from Specifications*. Prentice-Hall, second edition, 1994. web.comlab.ox.ac.uk/oucl/publications/books/PfS/.
 - 25 C.C. Morgan, A.K. McIver, and K. Seidel. Probabilistic predicate transformers. *ACM Trans Prog Lang Sys*, 18(3):325–53, May 1996. [doi.acm.org/10.1145/229542.229547](https://doi.org/10.1145/229542.229547).
 - 26 Geoffrey Smith. On the foundations of quantitative information flow. In Luca de Alfaro, editor, *Proc. 12th International Conference on Foundations of Software Science and Computational Structures (FoSSaCS '09)*, volume 5504 of *Lecture Notes in Computer Science*, pages 288–302, 2009.
 - 27 Geoffrey Smith. Connections between g -leakage and the Dalenius desideratum, December 2014. Presentation at Princess workshop meeting, Paris. www.lix.polytechnique.fr/comete/Projects/Princess/workshops/quantitative-information-3/.
 - 28 Colin D. Walter. MIST: an efficient, randomized exponentiation algorithm for resisting power analysis. In *Topics in Cryptology - CT-RSA 2002, The Cryptographer's Track at the RSA Conference, 2002, San Jose, CA, USA, February 18-22, 2002, Proceedings*, pages 53–66, 2002.

A

 Summary of terms and notations

These entries list in first-use order the points at which notation is introduced during the exposition: a detailed explanation of each is given there.

collateral leakage	Leakage from data to which a program does not refer.	p.1
\mathcal{X}	(Fig. 1) State space.	p.2
$[\cdot \cdots]$	(Fig. 1) Uniform distribution.	p.2
$:\in$	(Fig. 1) Assignment made according to a distribution.	p.2
leak	(Fig. 1) Emit information from a program that an adversary can observe; causes no change in the state of the program.	p.2
x^+, x^-	(Fig. 1) Next, previous letter in a list.	p.2
Lax, Strict	(Fig. 1) Example password programs..	p.2
collateral variable	variable affected by collateral leakage.	p.3
abstract channel	Denotation of channel that abstracts from leaked values.	p.3
abstract HMM	Denotation of HMM that abstracts from leaked values.	p.3
$\mathcal{X} \rightarrow \mathcal{Y}$	Row-stochastic matrix.	p.3
$\mathcal{X} \leftarrow \mathcal{Y}$	Column-stochastic (co-stochastic) matrix.	p.3
$C_{x,y}, C_{x,-}, C_{-,y}$	Element, row, column of a matrix.	p.3
\mathbb{D}	Discrete-distribution- or probability-measure type-constructor.	p.4
$\pi \rangle C$	Joint distribution on $\mathcal{Y} \times \mathcal{X}$ formed by taking prior π as input to channel C .	p.4
[joint distribution]	The hyper formed from a joint distribution.	p.4
hyper-distribution	Type \mathbb{D}^2 ; distributions of distributions.	p.4
$\mathbb{D}\mathcal{X} \rightarrow \mathbb{D}^2\mathcal{X}$	Type of abstract channel.	p.4
vulnerability	Quantitative measure of secret's value.	p.4
gain function	Used to construct vulnerability functions.	p.4
\mathcal{W}	Set from which gain-function's choices are taken.	p.4
$\mathcal{E}_\pi f$	Expected value of function f on distribution π .	p.4
$g: \mathbb{G}\mathcal{X}$	Typical gain function and its type.	p.4
g_{id}	The identity g -function, with $\mathcal{W}=\mathcal{X}$, giving $V_{g_{id}}$ as Bayes vulnerability.	p.4
$V_g: \mathbb{D}\mathcal{X} \rightarrow \mathbb{R}$	Vulnerability V_g induced by a gain function g , and its type.	p.4
Bayes Vulnerability	Greatest probability of guessing a secret on the first try.	p.5
Bayes Risk	One minus the Bayes Vulnerability.	p.5
\mathcal{L}_g etc.	Multiplicative g -leakage, and various capacities.	p.5
dot \cdot	Matrix multiplication.	p.5
$\tilde{\pi}, \tilde{\Pi}$	left-marginal, right channel of joint distribution	p.5
$(C:M)$	HMM -step, a "primitive" HMM with channel C and markov M .	p.6
markov (lc.)	The Markov portion of an HMM -step.	p.6
semicolon ;	Sequential composition of HMM 's as matrices.	p.6
chn	The effective channel of an HMM .	p.6
\parallel	Parallel composition of channels.	p.7
CCap	Collateral-capacity estimate.	p.7
$\Upsilon_{\mathcal{X}}$	The uniform distribution over \mathcal{X} .	p.7
extremal leakage	An upper bound on leakage.	p.7
$[A^{\oplus 1/2}, B^{\oplus 1/4}, C^{\oplus 1/4}]$	Discrete distribution with explicit probabilities.	p.9
wprob	with probability	p.9

\sqsubseteq	Security order on hypers.	p.9
$\llbracket \cdot \rrbracket$	Semantic function taking <i>HMM</i> matrices to abstract <i>HMM</i> 's.	p.10
$\pi \rangle H$	Joint distribution on $\mathcal{X} \times \mathcal{Y} \times \mathcal{X}'$ formed by taking prior π as input to hyper H .	p.10
$\mathbb{D}\mathcal{X} \rightarrow \mathbb{D}^2\mathcal{X}^2$	Type of abstract <i>HMM</i> , sometimes written $\mathbb{D}\mathcal{X} \rightarrow \mathbb{D}^2(\mathcal{X} \times \mathcal{X}')$, with $\mathcal{X} = \mathcal{X}'$, to emphasise the different roles (initial, final) of the two \mathcal{X} 's.	p.10
abstract <i>HMM</i>	Function in $\mathbb{D}\mathcal{X} \rightarrow \mathbb{D}^2(\mathcal{X} \times \mathcal{X}')$ that is $\llbracket H \rrbracket$ for some $H: \mathcal{X} \rightarrow \mathcal{Y} \times \mathcal{X}'$.	p.10
$h^{\mathcal{Z}}$	Utility function for extending collateral-aware <i>HMM</i> 's.	p.10
(bold) semicolon ;	Sequential composition of abstract <i>HMM</i> 's.	p.10
$h^{\times \mathcal{Z}}$	Collateral extension of abstract <i>HMM</i> to collateral variable \mathcal{Z} .	p.10
ζ	“Duplicator” utility function used in the extension $h^{\times \mathcal{Z}}$.	p.10
$\text{nc}_{\mathcal{X}}$	Degenerate channel on \mathcal{X} that leaks nothing.	p.25
$\text{id}_{\mathcal{X}}$	Degenerate Markov transition on \mathcal{X} that updates nothing.	p.25
pure channel	An <i>HMM</i> that (might) leak, but does not update the state.	p.25
pure markov	An <i>HMM</i> that (might) update the state, but does not leak.	p.25
$(:M)$	Abbreviation for $(\text{nc}:M)$, a pure markov.	p.25
$(C:)$	Abbreviation for $(C:\text{id})$, a pure channel.	p.25
\circ	The sole element of the unit type $\{\circ\}$ emitted by a degenerate channel.	p.18

// Password X is initially uniformly distributed over $\mathcal{X} = \{A, B, C\}$.

"Lax" user	
$X \leftarrow [A, B, C]$	*
$\text{leak } [X^+, X^-]$	†

"Strict" user	
$X \leftarrow [X^+, X^-]$	§
$\text{leak } [X^+, X^-]$	†

* [...] is the uniform distribution over ... ; and $X \leftarrow$ assigns to X from a distribution.

§ X^+ is the letter following X , and X^- the preceding.

† $\text{leak } [X^+, X^-]$ makes a (fair) choice secretly between X^+ or X^- , then emits the value somehow: it does not however indicate whether it chose X^+ or X^- .

Note that the X referred to in the leak statement is the updated value, after the $X \leftarrow$.

Lax may choose any new password, uniformly, including his current; but Strict must *change* his password, again uniformly. In both cases the distribution of the new X is again uniform: for Lax it is independent of X 's initial value; but for Strict is is correlated.

Both users, in the second statement, suffer an "over the shoulder" attack against the new password.

■ **Figure 4** Updating a password (repeat of Fig. 1)

B Detailed calculation of hypers for Fig. 1

[from §1, §5.1]

B.1 Elementary models

For convenience we repeat the programs of Fig. 1, in Fig. 4 here.

Variable X in the program text ranges over the state space $\mathcal{X} = \{A, B, C\}$. We begin by modelling its components in the simplest mathematical way, i.e. conventionally with probabilistic updates being Markov matrices in $\mathcal{X} \rightarrow \mathcal{X}$ and leaks being channel matrices in $\mathcal{X} \rightarrow \mathcal{Y}$, where in this simple example $\mathcal{Y} = \mathcal{X}$. We then map those different descriptions into the same framework of *HMM*-steps, their types becoming $\mathcal{X} \rightarrow \mathcal{Y} \times \mathcal{X}$, for varying \mathcal{Y} (sometimes \mathcal{X} itself, and sometimes the unit type $\{\circ\}$); and finally we use their sequential compositions to get the *HMM* descriptions of the two programs Lax and Strict.

The first statement $X \leftarrow [A, B, C]$ of Lax corresponds to a Markov matrix in $\mathcal{X} \rightarrow \mathcal{X}$ with the input at left and output at top:

$$M^{L1} : \begin{array}{c} \begin{array}{ccc} & A & B & C \\ \begin{array}{c} A \\ B \\ C \end{array} & \begin{pmatrix} 1/3 & 1/3 & 1/3 \\ 1/3 & 1/3 & 1/3 \\ 1/3 & 1/3 & 1/3 \end{pmatrix} \end{array} & \begin{array}{l} \text{For Lax the input is ig-} \\ \text{nored; the output is a} \\ \text{uniform choice over } \mathcal{X}. \end{array} \end{array}$$

However the first statement $X \leftarrow [X^+, X^-]$ of Strict corresponds this Markov matrix instead:

$$M^{S1} : \begin{array}{c} \begin{array}{ccc} & A & B & C \\ \begin{array}{c} A \\ B \\ C \end{array} & \begin{pmatrix} 0 & 1/2 & 1/2 \\ 1/2 & 0 & 1/2 \\ 1/2 & 1/2 & 0 \end{pmatrix} \end{array} & \begin{array}{l} \text{For Strict the output is} \\ \text{a uniform choice over} \\ \text{anything but the input.} \end{array} \end{array}$$

The second statement $X \leftarrow [A, B, C]$ in both cases corresponds to a channel matrix in $\mathcal{X} \rightarrow \mathcal{Y}$, where in fact $\mathcal{Y} = \mathcal{X}$ because the observables are of the same type as the state:

$$C^2 : \begin{array}{c} \begin{array}{ccc} & A & B & C \\ \begin{array}{c} A \\ B \\ C \end{array} & \begin{pmatrix} 0 & 1/2 & 1/2 \\ 1/2 & 0 & 1/2 \\ 1/2 & 1/2 & 0 \end{pmatrix} \end{array} & \begin{array}{l} \text{The over-the-shoulder leak} \\ \text{is uniformly any value not} \\ \text{equal to the current state.} \end{array} \end{array}$$

B.2 Elementary models make *HMM* steps

To build the two programs from their components, we must convert those components to *HMM* steps: note that although the components are all matrices, Markov matrices and channels are converted in different ways. We start with the conversion of the markovs: converted from $\mathcal{X} \rightarrow \mathcal{X}$ to the *HMM*-type $\mathcal{X} \rightarrow \mathcal{Y} \times \mathcal{X}$ they become

$$\begin{array}{ccc} \circ A & \circ B & \circ C \\ \begin{array}{l} A \\ B \\ C \end{array} \begin{pmatrix} 1/3 & 1/3 & 1/3 \\ 1/3 & 1/3 & 1/3 \\ 1/3 & 1/3 & 1/3 \end{pmatrix} & \begin{array}{l} \circ A & \circ B & \circ C \\ \begin{array}{l} A \\ B \\ C \end{array} \begin{pmatrix} 0 & 1/2 & 1/2 \\ 1/2 & 0 & 1/2 \\ 1/2 & 1/2 & 0 \end{pmatrix} \end{array}$$

where e.g. the label $\circ A$ means the y, x' pair (\circ, A) — the observable is “unit”, the single element of the unit type $\{\circ\}$, and the new state is A because, these markovs’ being leak-free, they (in effect) have no observables (equivalently, one always observes the same thing, the \circ , so no information flows).¹⁵

Conversion of the channel(s) does however leak information, giving in both cases¹⁶

$$\begin{array}{ccccccccc} AA & AB & AC & BA & BB & BC & CA & CB & CC \\ \begin{array}{l} A \\ B \\ C \end{array} \begin{pmatrix} 0 & 0 & 0 & 1/2 & 0 & 0 & 1/2 & 0 & 0 \\ 0 & 1/2 & 0 & 0 & 0 & 0 & 0 & 1/2 & 0 \\ 0 & 0 & 1/2 & 0 & 0 & 1/2 & 0 & 0 & 0 \end{pmatrix} & \begin{array}{l} \text{The columns are labelled with pairs } (y, x'), \\ \text{so that e.g. from input } A \text{ there are two} \\ \text{possible results: leak } B \text{ and stay in state } A, \\ \text{or leak } C \text{ and stay in state } A. \end{array} \end{array}$$

B.3 Sequential composition of *HMM*-steps

To get the *HMM* representation of the two programs, we use the definition of *HMM* sequential composition (2) in §3.1, which we repeat here for convenience:

$$(H^1; H^2)_{x, (y^1, y^2), x'} = \sum_{x''} H^1_{x, y^1, x''} H^2_{x'', y^2, x'} \quad \text{recalling (2)} \quad (4)$$

We have two versions of H^1 , one for each of **Lax** and **Strict**, and a single version of H^2 that applies to both. Type \mathcal{Y}^1 is the unit $\{\circ\}$; and type \mathcal{Y}^2 is the same as \mathcal{X} . Our composite *HMM* (no longer simply an *HMM*-step) will have observables’ type $\mathcal{Y} = \mathcal{X}$, because we can drop the unit type;¹⁷ and its output type is still \mathcal{X} . For **Lax** the result is this *HMM*, obtained by applying (2) to **Lax**’s H^1 and the common H^2 :

$$\begin{array}{ccccccccc} AA & AB & AC & BA & BB & BC & CA & CB & CC \\ \begin{array}{l} A \\ B \\ C \end{array} \begin{pmatrix} 0 & 1/6 & 1/6 & 1/6 & 0 & 1/6 & 1/6 & 1/6 & 0 \\ 0 & 1/6 & 1/6 & 1/6 & 0 & 1/6 & 1/6 & 1/6 & 0 \\ 0 & 1/6 & 1/6 & 1/6 & 0 & 1/6 & 1/6 & 1/6 & 0 \end{pmatrix} & \begin{array}{l} \text{For } \mathbf{Lax} \text{ the rows are identical, because the} \\ \text{input is ignored; the result is uniformly dis-} \\ \text{tributed over outcomes in which the leaked} \\ \text{value and the new state are different.} \end{array} \end{array}$$

For **Strict** the result is this *HMM*, this time using **Strict**’s H^1 (and again the common H^2):

$$\begin{array}{ccccccccc} AA & AB & AC & BA & BB & BC & CA & CB & CC \\ \begin{array}{l} A \\ B \\ C \end{array} \begin{pmatrix} 0 & 1/4 & 1/4 & 0 & 0 & 1/4 & 0 & 1/4 & 0 \\ 0 & 0 & 1/4 & 1/4 & 0 & 1/4 & 1/4 & 0 & 0 \\ 0 & 1/4 & 0 & 1/4 & 0 & 0 & 1/4 & 1/4 & 0 \end{pmatrix} & \begin{array}{l} \text{For } \mathbf{Strict} \text{ the rows are } \textit{not} \text{ identical, be-} \\ \text{cause the first } \mathbf{HMM}\text{-step updates the state} \\ \text{in a way dependent on its incoming value.} \end{array} \end{array}$$

¹⁵ In the notation of App. E they would be $(:M^{L1})$ and $(:M^{S1})$, both of *HMM*-type $\mathcal{X} \rightarrow \{\circ\} \times \mathcal{X}$.

¹⁶ In the notation of App. E this would be (C^2) of *HMM*-type $\mathcal{X} \rightarrow \mathcal{Y} \times \mathcal{X}$ where, recall in fact $\mathcal{Y} = \mathcal{X}$.

¹⁷ The Cartesian product $\{\circ\} \times \mathcal{Y}$ is isomorphic to \mathcal{Y} alone: we have *HMM*-composed $\mathcal{X} \rightarrow \{\circ\} \times \mathcal{X}$ and $\mathcal{X} \rightarrow \mathcal{Y} \times \mathcal{X}$ to get $\mathcal{X} \rightarrow (\{\circ\} \times \mathcal{Y}) \times \mathcal{X}$, isomorphically $\mathcal{X} \rightarrow \mathcal{Y} \times \mathcal{X}$.

B.4 The semantic view: hypers for Lax

We are now interested in the hypers produced by each of the two *HMM*'s just above when they are applied to the uniform prior $[A, B, C]$ that assigns probability $1/3$ to each of A, B, C , an assumption we have made for this example). We do **Lax** first: the joint distribution in $\mathbb{D}(\mathcal{X} \times \mathcal{Y} \times \mathcal{X}')$, where we are now referring to the final state with a prime, is obtained by multiplying each row of the *HMM* by the corresponding prior probability (in this case $1/3$ for all of them), giving

$$\begin{array}{c} \text{AA} \quad \text{AB} \quad \text{AC} \quad \text{BA} \quad \text{BB} \quad \text{BC} \quad \text{CA} \quad \text{CB} \quad \text{CC} \\ \begin{array}{l} \text{A} \\ \text{B} \\ \text{C} \end{array} \left(\begin{array}{ccccccccc} 0 & 1/18 & 1/18 & 1/18 & 0 & 1/18 & 1/18 & 1/18 & 0 \\ 0 & 1/18 & 1/18 & 1/18 & 0 & 1/18 & 1/18 & 1/18 & 0 \\ 0 & 1/18 & 1/18 & 1/18 & 0 & 1/18 & 1/18 & 1/18 & 0 \end{array} \right) \end{array} \quad \begin{array}{l} \text{To convert this to a hyper on } \mathcal{X} \times \mathcal{X}', \\ \text{we abstract from } \mathcal{Y}, \text{ the first component} \\ \text{of each column.} \end{array}$$

Because of the symmetry, to abstract from \mathcal{Y} we can concentrate on just one observed value, say **A**. In that case, the posterior in $\mathbb{D}(\mathcal{X} \times \mathcal{X}')$ is $[[AB, AC], [BB, BC], [CB, CC]]$, where the first component is the row label x from just above, and the second component is the second component of the column label x' — the observation $y=A$ has been abstracted. We can write that more concisely in this case as the product $[A, B, C] \times [B, C]$ of two (uniform) distributions, one in $\mathbb{D}\mathcal{X}$ and the other in $\mathbb{D}\mathcal{X}'$.

For observation **B**, analogously we get $[A, C] \times [A, C]$; for observation **C** we get $[A, C] \times [A, B]$. Since the observations (in this case) have equal (marginal) probabilities, the overall hyper in $\mathbb{D}^2(\mathcal{X} \times \mathcal{X}')$ is uniform: it is

$$[[A, B, C] \times [B, C], [A, B, C] \times [A, C], [A, B, C] \times [A, B]] . \quad (5)$$

If we project (5) onto its second component we get a hyper in $\mathbb{D}^2\mathcal{X}'$ on the final state alone: it is $[[B, C], [A, C], [A, B]]$, so that for **Lax** we will wprob $1/3$ believe a-posteriori that x' is uniformly either **B** or **C**, when we observed **A**; for observations **B, C** resp. our posteriors in $\mathbb{D}\mathcal{X}'$ alone are $[A, C]$ and $[A, B]$ resp.

If however we project onto the first component, we get posterior $[A, B, C]$ in all three cases: we have learned nothing about the *initial* value of x .

B.5 The semantic view: hypers for Strict

For **Strict**, sensitive to the initial state, the results are different. First, the joint distribution in $\mathbb{D}(\mathcal{X} \times \mathcal{Y} \times \mathcal{X}')$ is again obtained by multiplying the uniform prior through, giving now

$$\begin{array}{c} \text{AA} \quad \text{AB} \quad \text{AC} \quad \text{BA} \quad \text{BB} \quad \text{BC} \quad \text{CA} \quad \text{CB} \quad \text{CC} \\ \begin{array}{l} \text{A} \\ \text{B} \\ \text{C} \end{array} \left(\begin{array}{ccccccccc} 0 & 1/12 & 1/12 & 0 & 0 & 1/12 & 0 & 1/12 & 0 \\ 0 & 0 & 1/12 & 1/12 & 0 & 1/12 & 1/12 & 0 & 0 \\ 0 & 1/12 & 0 & 1/12 & 0 & 0 & 1/12 & 1/12 & 0 \end{array} \right) \end{array}$$

Again we concentrate on just one observation **A**, in which case, the posterior in $\mathbb{D}(\mathcal{X} \times \mathcal{X}')$ is $[AB, AC, BC, CB]$.¹⁸ For observation **B**, analogously we get $[AC, BA, BC, CA]$; for observation **C** we get $[AB, BA, CA, CB]$. Since the observations (again) have equal probabilities, the overall hyper in $\mathbb{D}^2(\mathcal{X} \times \mathcal{X}')$ is again uniform:

$$[[AB, AC, BC, CB], [AC, BA, BC, CA], [AB, BA, CA, CB]] . \quad (6)$$

¹⁸Take each of the non-zero entries under a column with left component **A**; retain only x and x' , dropping the **A**; and, since all the probabilities are the same, form the uniform distribution on that.

Now if we project this hyper, for *Strict*, onto its second component (the final state), we get the same hyper in $\mathbb{D}^2\mathcal{X}'$ as for *Lax* just above. But if we project onto the first component (the initial state), this time we get posterior

$$[[A, A, B, C], [A, B, B, C], [A, B, C, C]], \quad (7)$$

which we constructed just by erasing the second component of each pair in (6). The “uniform distribution” with repeated components in the list gives proportionally more probability to the repetitions: written in the conventional way, for (7) we would get

$$[[A^{\otimes 1/2}, B^{\otimes 1/4}, C^{\otimes 1/4}], [A^{\otimes 1/4}, B^{\otimes 1/2}, C^{\otimes 1/4}], [A^{\otimes 1/4}, B^{\otimes 1/4}, C^{\otimes 1/2}]], \quad (8)$$

which is indeed the same as (3) in §5.1 above.

B.6 The semantics of composition, and relative security

The above constructions started from Markov- and channel matrices, converted them into *HMM*-step matrices, and then sequentially composed those matrices according to (4). The semantic alternative is to take the denotations at the very beginning, i.e. the abstract *HMM*’s corresponding to the individual program fragments, and then to compose them according to the *semantic* definition at Def. 8. The essential property $\llbracket H^1; H^2 \rrbracket = \llbracket H^1 \rrbracket; \llbracket H^2 \rrbracket$ referred to there ensures that the outcome would have been the same. Thus, either way, we can summarise all the above by saying that *Lax*, as an abstract *HMM*, takes prior $[A, B, C]$ to hyper (5) which, when projected onto the initial component \mathcal{X} is $[[A, B, C]]$ — showing that indeed *Lax* releases nothing about the initial state. The same procedure with *Strict* however gives the initial-state hyper (8) which is *strictly less secure* than $[[A, B, C]]$ in the hyper-order \sqsubseteq . This can be demonstrated even with Bayes Vulnerability $V_{g_{id}}$, which is $1/3$ for $[[A, B, C]]$ but $1/2$ for (8) — but, more significantly, that strict-less-secure relation means that *for any* g applying V_g to $[[A, B, C]]$ will give no worse (i.e. no higher) vulnerability than applying it to (8) or, equivalently, to (3).

C

 Proofs for §3

Proof for Lem. 3

[from §3.2]

Let H be an HMM and $(C:M)$ an HMM -step. Then

$$\begin{aligned} \text{chn.}(C:M) &= C \\ \text{chn.}((C:M); H) &= C \parallel (M \cdot \text{chn.}H) \end{aligned}$$

where in general $(C^1 \parallel C^2)_{x, (y^1, y^2)} = C^1_{x, y^1} C^2_{x, y^2}$ is parallel composition of channels. The M cannot be discarded, since it affects the prior of the “tail” H of the sequential composition.

Proof. We prove each equality independently:

$$\begin{aligned} & (\text{chn.}(C:M))_{x, y} \\ = & \sum_{x'} (C:M)_{x, y, x'} && \text{“Def. chn”} \\ = & \sum_{x'} C_{x, y} M_{x, x'} && \text{“Def. HMM-matrix”} \\ = & C_{x, y} && \text{“}M \text{ is stochastic”} \\ \\ & (\text{chn.}((C:M); H))_{x, (y_1, \dots, y_n)} \\ = & \sum_{x'} ((C:M); H)_{x, (y_1, \dots, y_n), x'} && \text{“Def. chn”} \\ = & \sum_{x'} \sum_{x''} (C:M)_{x, y_1, x''} H_{x'', (y_2, \dots, y_n), x'} && \text{“Def. (;)”} \\ = & \sum_{x'} \sum_{x''} C_{x, y_1} M_{x, x''} H_{x'', (y_2, \dots, y_n), x'} && \text{“Def. HMM-matrix”} \\ = & C_{x, y_1} \sum_{x''} M_{x, x''} \sum_{x'} H_{x'', (y_2, \dots, y_n), x'} && \text{“Move stuff around”} \\ = & C_{x, y_1} \sum_{x''} M_{x, x''} (\text{chn.}H)_{x'', (y_2, \dots, y_n)} && \text{“Def. chn”} \\ = & C_{x, y_1} (M \cdot \text{chn.}H)_{x, (y_2, \dots, y_n)} && \text{“Matrix multiplication”} \\ = & (C; (M \cdot \text{chn.}H))_{x, (y_1, y_2, \dots, y_n)} && \text{“Parallel composition”} \end{aligned}$$

◀

Proof for Lem. 4

[from §3.2]

For any H let $\text{CCap.}H$ be defined

$$\begin{aligned} \text{CCap.}(C:M) &= \mathcal{L}_\forall(\forall, C) && \text{if } H=(C:M) \\ \text{CCap.}((C:M); H) &= \mathcal{L}_\forall(\forall, C) + \min(\mathcal{L}_\forall(\forall, M), \text{CCap.}H) . && \text{if } H=(C:M); H' \end{aligned}$$

Then $\mathcal{L}_\forall(\forall, \text{chn.}H) \leq \text{CCap.}H$ where we interpret the stochastic matrix M (rhs) as a channel.

Proof. We use induction in the number of HMM -steps of H .

Base case: H has only one step.

$$\begin{aligned} & \mathcal{L}_\forall(\forall, \text{chn.}(C:M)) \\ = & \mathcal{L}_\forall(\forall, C) && \text{“Def. chn”} \\ = & \text{CCap.}(C:M) && \text{“Def. CCap”} \end{aligned}$$

Inductive step: assume the lemma is true for HMM ’s composed of n steps. Let $H = (C:M); H'$, where H' has n steps.

$$\begin{aligned}
& \mathcal{L}_\forall(\forall, \text{chn}((C:M); H')) \\
= & \mathcal{L}_\forall(\forall, C \parallel (M \cdot \text{chn}.H')) && \text{"Def. chn"} \\
= & \mathcal{L}_{g_{id}}(\Upsilon_{\mathcal{X}}, C \parallel (M \cdot \text{chn}.H')) && \text{"[2, Thm. 5.1]"} \\
\leq & \mathcal{L}_{g_{id}}(\Upsilon_{\mathcal{X}}, C) + \mathcal{L}_{g_{id}}(\Upsilon_{\mathcal{X}}, M \cdot \text{chn}.H') && \text{"[12, Cor. 7]"} \\
\leq & \mathcal{L}_{g_{id}}(\Upsilon_{\mathcal{X}}, C) + \min(\mathcal{L}_{g_{id}}(\Upsilon_{\mathcal{X}}, M), \mathcal{L}_{g_{id}}(\Upsilon_{\mathcal{X}}, \text{chn}.H')) && \text{"[12, Thm. 6]"} \\
= & \mathcal{L}_{g_{id}}(\Upsilon_{\mathcal{X}}, C) + \min(\mathcal{L}_{g_{id}}(\Upsilon_{\mathcal{X}}, M), \mathcal{L}_\forall(\forall, \text{chn}.H')) && \text{"[2, Thm. 5.1]"} \\
\leq & \mathcal{L}_{g_{id}}(\Upsilon_{\mathcal{X}}, C) + \min(\mathcal{L}_{g_{id}}(\Upsilon_{\mathcal{X}}, M), \text{CCap}.H') && \text{"Inductive hypothesis"} \\
\leq & \mathcal{L}_\forall(\forall, C) + \min(\mathcal{L}_\forall(\forall, M), \text{CCap}.H') && \text{"monotonicity"} \\
= & \text{CCap}((C:M); H') && \text{"Def. CCap"}
\end{aligned}$$

◀

Example showing that Lem. 4 can be conservative

[from §3.2]

Consider the following program fragment.

```
// Prog:  xs is initialised uniformly at random.
xs := xs  $\oplus_{1/3}$   $\neg$ xs ;
leak xs[0]  $\oplus_{1/2}$  xs[1] ;
xs := xs  $\oplus_{1/2}$   $\neg$ xs
```

Hidden variable **xs** is set initially uniformly at random over the four possible arrays of length two having 0,1 entries. Next the array values are all flipped (probability 2/3) or all left alone (probability 1/3); afterwards the value of either the first or second bit is leaked, but only the value is observed. Finally **xs** is updated again.

We model this as three *HMM* basic steps : $H := (:M_1); (C_2); (:M_2)$, where M_1 corresponds to the first update of **xs**, C_2 corresponds to the **leak** statement and M_2 the last update of **xs**. But from App. E we can combine the second and third, writing it as just the two steps $(:M_1); (C_2; M_2)$. Applying Lem. 4 we find

$$\begin{aligned}
& \text{CCap}.H \\
= & \mathcal{L}_{g_{id}}(\Upsilon_{\mathcal{X}}, I) + \min(\mathcal{L}_{g_{id}}(\Upsilon_{\mathcal{X}}, M_1), \text{CCap}((C_2; M_2))) \\
= & \min(\mathcal{L}_{g_{id}}(\Upsilon_{\mathcal{X}}, M_1), \mathcal{L}_{g_{id}}(\Upsilon_{\mathcal{X}}, C_2)) . && \text{"Identity channel leaks nothing; Lem. 4"}
\end{aligned}$$

We compute the values of the relevant matrices as follows:

$$M_1 := \begin{pmatrix} 1/3 & 0 & 0 & 2/3 \\ 0 & 1/3 & 2/3 & 0 \\ 0 & 2/3 & 1/3 & 0 \\ 2/3 & 0 & 0 & 1/3 \end{pmatrix} \quad C_2 := \begin{pmatrix} 1 & 0 \\ 1/2 & 1/2 \\ 1/2 & 1/2 \\ 0 & 1 \end{pmatrix}$$

Finally we compute the leakages:

$$\begin{aligned}
\mathcal{L}_{g_{id}}(\Upsilon_{\mathcal{X}}, M_1) &= \lg((1/6 + 1/6 + 1/6 + 1/6) * 4) = \lg(8/3) = 1.415 , \\
\mathcal{L}_{g_{id}}(\Upsilon_{\mathcal{X}}, C_2) &= \lg((1/4 + 1/4) * 4) = \lg(2) = 1 .
\end{aligned}$$

Hence $\text{CCap}.H = 1$.

However we can get an exact computation of leakage by calculating $\text{chn}.H$ exactly via Lem. 3, which gives $\mathcal{L}_\forall(\forall, \text{chn}.H) = \mathcal{L}_{g_{id}}(\Upsilon_{\mathcal{X}}, M_1 \cdot C_2)$ where

$$M_1 \cdot C_2 := \begin{pmatrix} 1/3 & 2/3 \\ 1/2 & 1/2 \\ 1/2 & 1/2 \\ 2/3 & 1/3 \end{pmatrix}$$

So finally we have $\mathcal{L}_V(\forall, \text{chn}.H) = \lg((1/6 + 1/6)4) = \lg(4/3) = 0.415$, which is considerably less than $\text{CCap}.H$.

Proof for Thm. 5

[from §3.2]

Given H and $\Pi: \mathbb{D}(\mathcal{Z} \times \mathcal{X})$ with $\tilde{\pi}, \bar{\pi}$ resp. the marginals of Π on \mathcal{Z}, \mathcal{X} , define conditional $\bar{\Pi}$ as in §2.4, there exists $\hat{g}: \mathbb{G}_{\mathcal{Z}} \mathcal{Z}$ and $\hat{g}^\Pi: \mathbb{G}_{\mathcal{Z}} \mathcal{X}$ such that

$$\mathcal{L}_V(\tilde{\pi}, \bar{\Pi} \cdot \text{chn}.H) = \mathcal{L}_{\hat{g}}(\tilde{\pi}, \bar{\Pi} \cdot \text{chn}.H) = \mathcal{L}_{\hat{g}^\Pi}(\tilde{\pi}, \text{chn}.H) .$$

Proof. Let $a := \min_{z: \mathcal{X}, \tilde{\pi}_z > 0} \tilde{\pi}_z$, that is, a is the minimum non-zero probability in $\tilde{\pi}$. The first equality follows from [1][Theorem 10], with $\hat{g}.z'.z = a/\tilde{\pi}_z$, whenever $z' = z$ and $\tilde{\pi}_z > 0$, and is zero otherwise.

We prove the second equality as follows. Let $C := \text{chn}.H$ and let $\bar{\Pi}$ in $\mathcal{X} \rightarrow \mathcal{Z}$ be such that $\Pi_{z,x} = \tilde{\pi}_x \bar{\Pi}_{x,z}$. Next, given $g: \mathbb{G}_{\mathcal{W}} \mathcal{Z}$, define $g^\Pi: \mathbb{G}_{\mathcal{W}} \mathcal{X}$ by $g^\Pi.w.x = \sum_{z: \mathcal{Z}} \bar{\Pi}_{x,z} g.w.z$.

We now reason as follows that $V_g[\tilde{\pi} \bar{\Pi} \cdot C] = V_{g^\Pi}[\tilde{\pi} \cdot C]$.

$$\begin{aligned} & V_g[\tilde{\pi} \bar{\Pi} \cdot C] \\ = & \sum_{y: \mathcal{Y}} \max_{w: \mathcal{W}} \sum_{z: \mathcal{Z}} \tilde{\pi}_z (\bar{\Pi} \cdot C)_{z,y} g.w.z && \text{“Def. } V_g \text{”} \\ = & \sum_{y: \mathcal{Y}} \max_{w: \mathcal{W}} \sum_{z: \mathcal{Z}} \tilde{\pi}_z (\sum_{x: \mathcal{X}} \bar{\Pi}_{z,x} C_{x,y}) g.w.z && \text{“cascading”} \\ = & \sum_{y: \mathcal{Y}} \max_{w: \mathcal{W}} \sum_{z: \mathcal{Z}, x: \mathcal{X}} (\tilde{\pi}_z \bar{\Pi}_{z,x}) C_{x,y} g.w.z && \text{“move stuff”} \\ = & \sum_{y: \mathcal{Y}} \max_{w: \mathcal{W}} \sum_{z: \mathcal{Z}, x: \mathcal{X}} (\tilde{\pi}_x \bar{\Pi}_{x,z}) C_{x,y} g.w.z && \text{“}\tilde{\pi}_z \bar{\Pi}_{z,x} = \Pi_{z,x} = \tilde{\pi}_x \bar{\Pi}_{x,z} \text{”} \\ = & \sum_{y: \mathcal{Y}} \max_{w: \mathcal{W}} \sum_{x: \mathcal{X}} \tilde{\pi}_x C_{x,y} (\sum_{z: \mathcal{Z}} \bar{\Pi}_{x,z} g.w.z) && \text{“move stuff”} \\ = & \sum_{y: \mathcal{Y}} \max_{w: \mathcal{W}} \sum_{x: \mathcal{X}} \tilde{\pi}_x C_{x,y} g^\Pi.w.x && \text{“Def. } g^\Pi \text{”} \\ = & V_{g^\Pi}[\tilde{\pi} \cdot C] . && \text{“Def. } V_{g^\Pi} \text{”} \end{aligned}$$

The proof of $V_g[\tilde{\pi}] = V_{g^\Pi}[\tilde{\pi}]$ is similar, implying that $\mathcal{L}_{\hat{g}}(\tilde{\pi}, \bar{\Pi} \cdot \text{chn}.H) = \mathcal{L}_{\hat{g}^\Pi}(\tilde{\pi}, \text{chn}.H)$ as required. \blacktriangleleft

Proof for Thm. 6

[from §3.2]

Given H and Π as above, we have

$$\mathcal{L}_V(\tilde{\pi}, \bar{\Pi} \cdot \text{chn}.H) \leq \mathcal{L}_{g_{id}}(\Upsilon_{\mathcal{X}}, \text{chn}.H) ,$$

where $\tilde{\pi}, \bar{\Pi}$ are as defined in Thm. 5, and $\Upsilon_{\mathcal{X}}$ is the uniform probability distribution over \mathcal{X} .

Proof. It is enough to show the following:

$$\mathcal{L}_V(\tilde{\pi}, \bar{\Pi} \cdot \text{chn}.H) \leq \mathcal{L}_V(\tilde{\pi}, \text{chn}.H) = \mathcal{L}_{g_{id}}(\Upsilon_{\mathcal{X}}, \text{chn}.H)$$

The first inequality follows from the right capacity bound from [1, Cor 23]; the equality then follows from [1, Thm 10]. \blacktriangleleft

```

// B for base, the cleartext; E for exponent, the key: precondition is B,E >= 0,0 .
// P for power, the ciphertext.
P:= 1
while E!=0 // Invariant is  $P \cdot (B^E) = b^e$ , where  $b,e$  are initial values of B,E .
  D:∈ [2,3,5] // D for divisor; uniform choice from {2,3,5}.
  R:= E mod D; // R for remainder.
  if R!=0 then P:= P*B^R fi // Side-channel: is E divisible exactly by D ?
  B:= B^D // D is small: assume no side-channel here.
  E:= E div D // State update of E here. (No side-channel.)
end
// Now  $P=b^e$  and E=0: but what has an adversary learned about the initial  $e$  ?

```

Although our state comprises B,E,P,D,R we concentrate only on the secrecy of E. In particular, we are not trying to discover B or P in this case; and D,R are of no external significance afterwards anyway.

■ **Figure 5** Defence against side channel analysis in exponentiation: a repeat of Fig. 2

D An example: Obfuscated exponentiation

[from §4]

In this section give the details of the leakage calculations for the example from §4. For convenience we repeat the program at Fig. 5.

The well-known fast method of computing B^E from a base B exponent E is to employ a “divide and conquer” method to avoid the exponential time taken to perform E multiplications. But in the most straightforward implementations, side-channels can reveal whether the current value of the exponent was decremented or divided by two, effectively revealing the exponent’s bits one-by-one from least- to most significant. In cryptographic applications, this reveals the secret key [28].

A defence against that is to vary unpredictably the divisor that decreases the exponent on each iteration: not always 2, but say sometimes 3 and sometimes 5. This makes the algorithm (slightly) less efficient but, in compensation, much complicates the side-channel analysis that in the 2-case reveals everything. The code in Fig. 5 shows how: a fresh random divisor D is chosen from a fixed set \mathcal{D} , independently each time round the loop; in Fig. 5 we use $\mathcal{D} = \{2, 3, 5\}$. (The original, insecure algorithm uses $\mathcal{D}=\{2\}$.) Since a higher divisor offers more security, but takes longer to compute, the precise choice of \mathcal{D} can be thought of as selecting a trade-off between information leakage and performance.¹⁹

We assume that all variables are secret: for the example we assume the adversary seeks the exponent E only, that is the secret key, not e.g. a particular message. His side-channel is the if-statement: where he learns whether E is exactly divisible by D ; complicating his analysis however is that he does not know exactly what D is on each occasion.

We model this as an *HMM* as follows. We represent the state as a tuple (B, E, P, D, R) containing the values of each variable in the program. We can then model the statement $D: \in [2, 3, 5]$ as a Markov update M_D that uniformly assigns a fresh value to D while leaving the rest of the variables untouched. That is, it maps each tuple (B, E, P, D, R) to a tuple (B, E, P, D', R) with probability $1/|\mathcal{D}|$ for each $D' \in \mathcal{D}$. The assignment then following can be modeled as a deterministic Markov M_R , mapping (B, E, P, D, R) to $(B, E, P, D, E \bmod D)$. We

¹⁹Walter [28] gives a nice summary of various issues offering tradeoffs between leakage and performance. That algorithm is intended to provide more security by obfuscating further the pattern of if-branching. For simplicity we do not do that full analysis here.

assume pessimistically that the observer can learn whether or not the statement $P := P * B \wedge R$ was executed, which logically is equivalent to observing the evaluation of the guard $R! = 0$ of the `if` statement. We model this flow as a an *HMM* comprising a deterministic channel C and a Markov M_P . C outputs, for each state (B, E, P, D, R) , whether R is 0 or not. The (also deterministic) Markov part M_P updates the state (B, E, P, D, R) to (B, E, PB^R, D, R) only when $R \neq 0$, otherwise the state remains unchanged. Finally, the last two assignments correspond to deterministic Markov updates M_B and M_E , mapping a state (B, E, P, D, R) to (B^D, E, P, D, R) and $(B, E \div D, P, D, R)$ respectively.

We can express the program of Fig. 2 as an *HMM* built as the sequential composition

$$(:M_D); (:M_R); (C:M_P); (:M_B); (:M_E); \quad (:M_D); (:M_R); (C:M_P); \dots,$$

using the *HMM*-step notations from §3.1, whose good behaviour in representing our intentions is justified in App. E.

However, the exact number of iterations depends on the initial value of E . To simplify the analysis we will always assume a fixed number of iterations (equal to the minimum number of iterations required for the program to yield the right result in every case). For instance, if E can take values from 0 to 15 (4 bits), then the maximum number of iterations is 4: on each iteration, E is divided at least by 2, and therefore after 4 iterations the value of E is guaranteed to be 0. (“Extra iterations”, i.e. when E is already zero, have no effect.) Once we have an *HMM*, say H , we can calculate $\text{chn}.H$ using Lem. 3 and, once we have $\text{chn}.H$, we can calculate its min-leakage. From Thm. 6, that gives also Dalenius leakage, which in this case represents the leakage related to the initial value of E . Tbl. 1 shows the result of computing this leakage for different sets of divisors.

E Some notations and properties of *HMM* steps [used in App. D]

In §3.1 it was remarked that the single, uniform definition of $(;)$ for *HMM*’s “essentially” treats pure markovs and pure channels differently. That difference is possible because markovs and channels are encoded differently within *HMM*’s. First we give the details for classical *HMM*’s; then further below we address abstract *HMM*’s. The *degenerate* channel that leaks nothing about \mathcal{X} can be regarded as having an anonymous singleton set as its observations, a single-column matrix containing only 1’s. We write it $\text{nc}_{\mathcal{X}}$ for “no channel” on state-space \mathcal{X} ; usually we omit the \mathcal{X} . The degenerate markov that makes no change is the identity (matrix) on \mathcal{X} , which we write $\text{id}_{\mathcal{X}}$. We abbreviate $(\text{nc}:M)$ by $(:M)$ and $(C:\text{id})$ by $(C:)$, calling them (as *HMM*’s) *pure* channels and pure markovs, and we note these properties of composition:

1. $(C:); (:M) = (C:M)$, i.e. that $(C:M)$ is equivalent to “just C ” and then “just M ”.
2. $(C^1:); (C^2:) = (C^1 \parallel C^2:)$ where $(C^1 \parallel C^2)_{x, (y^1, y^2)}$ defined $C^1_{x, y^1} C^2_{x, y^2}$ is the *parallel composition* of its constituents, i.e. that “ C^1 and then C^2 ” is the same as C^1 and C^2 together because C^1 passes the state on, unchanged, to C^2 .
3. $(:M^1); (:M^2) = (:M^1 \cdot M^2)$, where $M^1 \cdot M^2$ is the (usual) Markov composition of its constituents, i.e. matrix multiplication.

The *HMM*-composition of pure channels and composition of pure markovs is what allows us to combine channels and markovs and then treat them with the same definition of sequential composition, as is done in App. D with one long (associative) sequential composition of the loop body’s constituents, then unfolded as explained there to eliminate the loop: their different purposes are automatically respected.

The essential property referred to in Def. 8 of §5.2 guarantees that this same respect for markovs' difference from channels is achieved in the semantic domain as well. Using [21, 22] this can be exploited to remain entirely in the semantic domain, including the loop explicitly without having to unfold it.

F

 Proofs and materials supporting §5

On “healthiness” of abstract *HMM*’s [from §5.2]

The denotations of our H ’s in $\mathcal{X} \rightarrow \mathcal{Y} \times \mathcal{X}$ are functions in $\mathbb{D}\mathcal{X} \rightarrow \mathbb{D}^2(\mathcal{X} \times \mathcal{X}')$; but not all functions in that space are $\llbracket H \rrbracket$ for some H . The ones that have been called abstract *HMM*’s, and they have properties that we use in the proofs below.

A common technique in semantics is to identify *characteristic* properties of denotations, which can then lead to further insights. They are expressed in semantic terms, as we did for example in [22, VIIA] for collateral-*unaware* abstract *HMM*’s, calling them “healthiness conditions”.²⁰ An advantage of that is the possible discovery of more general properties that secure computations’ (semantics) should have. Although we do not introduce collateral-aware healthiness here, it is a clear target for future work.

Proofs of lemmas and theorems [from §5.2]

Lemma 11 shows that h^Z in Fig. 3 is well defined whenever $h = \llbracket H \rrbracket$ for some H . The principal fact is that possible variation in the choice of the right-conditional Z , at left, does not affect the value produced by $\mathbb{D}(Z \cdot)$ at right.

► **Lemma 11.** *Let h be an abstract *HMM*; then h^Z from Fig. 3 is well defined.* [from §5.2]

Proof. Let $h = \llbracket H \rrbracket$ for some $H: \mathcal{X} \rightarrow \mathcal{Y} \times \mathcal{X}$, and let Π be a joint distribution in $\mathbb{D}(\mathcal{Z} \times \mathcal{X})$ with right marginal $\pi: \mathbb{D}\mathcal{X}$ and $Z: \mathcal{Z} \leftarrow \mathcal{X}$ a right-conditional so that $\Pi = Z \langle \pi$. If Z is unique, then the result is trivial because $h^Z \cdot \Pi$ is given by the composition of the bottom- and right-hand arrows acting on π in Fig. 3.

If however Z is not unique, it can only be at elements x where π_x is 0, the well known issue that the all-zero column $\Pi_{-,x}$ cannot be normalised: but equally well known is that “it doesn’t matter” because any arbitrary value chosen for $Z_{-,x}$ will be multiplied by 0 in $\Pi = Z \langle \pi$.

It doesn’t matter on the right, either, because a property of abstract h ’s inherited from their originating H ’s is that if $\pi.x = 0$ then the sub-vector $(\pi \rangle H)_{x,-,-}$ of the joint distribution $\pi \rangle H$ is all zero; that means, in turn, that in all inners δ of the hyper $[\pi \rangle H]$ derived from that joint distribution (by abstracting from y) the probability $\delta_{x,x'}$ for that x , and for for any x' , will be zero as well. Thus, again, variation of Z at that x does not matter.

In all cases,

$$h^Z \cdot \Pi = \mathbb{D}(Z \cdot) \cdot (h \cdot \pi) \tag{9}$$

holds whenever $\Pi = Z \langle \pi$. ◀

²⁰ A conspicuous example of this is Dijkstra’s healthiness conditions for predicate transformers on sequential programs: the principal one was conjunctivity. Originally [10] Dijkstra proved by structural induction over program texts that all program denotations were healthy; later researchers instead used healthiness to characterise a subspace of predicate transformers, and then proved that all programs’ denotations lay within it — the same proof, but a different point of view. That more semantic view allowed later the expansion of healthiness, to include e.g. miracles and angelic choice, which in turn lead to new programming-language features that could denote the members of the expanded set.

Proof for Thm. 10

[from §5.2]

Since $h^{\times \mathcal{Z}} = h^{\mathcal{Z}^2 \times \mathcal{X}} \circ \mathbb{D}\zeta$, it suffices to prove the following result.

► **Theorem 12.** *Let $h_{1,2}$ be two abstract HMM's. Then $h_1 \sqsubseteq h_2$ iff $h_1^{\mathcal{Z}} \sqsubseteq h_2^{\mathcal{Z}}$ for all extension \mathcal{Z} .*

Proof. Let us assume that $h_1 \sqsubseteq h_2$ and let $\Pi: \mathbb{D}(\mathcal{Z} \times \mathcal{X})$ with a right-conditional \mathcal{Z} and right marginal π , i.e. $\mathcal{Z} \langle \pi = \Pi$. Let $\Delta_i = h_i \cdot \pi$ for $i = 1, 2$. Recall from [20, Def. 6] that two hypers satisfy $\Delta_1 \sqsubseteq \Delta_2$ in $\mathbb{D}^2 \mathcal{X}^2$ iff there exists a super $\Gamma: \mathbb{D}^3 \mathcal{X}^2$ such that $\Delta_1 = \mu_{\mathbb{D} \mathcal{X}^2} \cdot \Gamma$ (the outer average) and $\Delta_2 = \mathbb{D} \mu_{\mathcal{X}^2} \cdot \Gamma$ (the inner average)²¹. We say that Γ *witnesses* the refinement $\Delta_1 \sqsubseteq \Delta_2$. Let us show that the super $\mathbb{D}^2(\mathcal{Z} \cdot) \cdot \Gamma: \mathbb{D}^3(\mathcal{Z} \times \mathcal{X})$ witnesses the refinement $\mathbb{D}(\mathcal{Z} \cdot) \cdot \Delta_1 \sqsubseteq \mathbb{D}(\mathcal{Z} \cdot) \cdot \Delta_2$ in $\mathbb{D}(\mathcal{Z} \times \mathcal{X})$, i.e.

$$\mathbb{D}(\mathcal{Z} \cdot) \cdot \Delta_1 = \mu_{\mathbb{D}(\mathcal{Z} \times \mathcal{X})} \cdot (\mathbb{D}^2(\mathcal{Z} \cdot) \cdot \Gamma) \quad \text{and} \quad \mathbb{D}(\mathcal{Z} \cdot) \cdot \Delta_2 = \mathbb{D} \mu_{\mathcal{Z} \times \mathcal{X}} \cdot (\mathbb{D}^2(\mathcal{Z} \cdot) \cdot \Gamma) .$$

Both equalities can be stated using the diagrams given in Fig. 6-(a,b). Let us show that these diagrams indeed commute.

$$\begin{array}{ccc}
 \begin{array}{ccc}
 \mathbb{D}^2 \mathcal{X}^2 & \xrightarrow{\mathbb{D}(\mathcal{Z} \cdot)} & \mathbb{D}^2(\mathcal{Z} \times \mathcal{X}) \\
 \uparrow \mu_{\mathcal{X}^2} & & \uparrow \mathbb{D} \mu_{(\mathcal{Z} \times \mathcal{X})} \\
 \mathbb{D}^3 \mathcal{X}^2 & \xrightarrow{\mathbb{D}^2(\mathcal{Z} \cdot)} & \mathbb{D}^3(\mathcal{Z} \times \mathcal{X})
 \end{array} &
 \begin{array}{ccc}
 \mathbb{D}^2 \mathcal{X}^2 & \xrightarrow{\mathbb{D}(\mathcal{Z} \cdot)} & \mathbb{D}^2(\mathcal{Z} \times \mathcal{X}) \\
 \uparrow \mu_{\mathcal{X}^2} & & \uparrow \mu_{\mathbb{D}(\mathcal{Z} \times \mathcal{X})} \\
 \mathbb{D}^3 \mathcal{X}^2 & \xrightarrow{\mathbb{D}^2(\mathcal{Z} \cdot)} & \mathbb{D}^3(\mathcal{Z} \times \mathcal{X})
 \end{array} &
 \begin{array}{ccc}
 \mathbb{D} \mathcal{X}^2 & \xrightarrow{(\mathcal{Z} \cdot)} & \mathbb{D}(\mathcal{Z} \times \mathcal{X}) \\
 \uparrow \mu_{\mathcal{X}^2} & & \uparrow \mu_{\mathcal{Z} \times \mathcal{X}} \\
 \mathbb{D}^2 \mathcal{X}^2 & \xrightarrow{\mathbb{D}(\mathcal{Z} \cdot)} & \mathbb{D}^2(\mathcal{Z} \times \mathcal{X})
 \end{array} \\
 \text{(a)} & \text{(b)} & \text{(c)}
 \end{array}$$

■ **Figure 6** The multiplication $(\mathcal{Z} \cdot)$ commutes with inner $(\mathbb{D} \mu)$ and outer $(\mu_{\mathbb{D}})$ averages.

For the diagram in Fig. 6-(a)(inner-average), it suffices to show Fig. 6-(c) commutes because \mathbb{D} -lifting preserves commutative diagrams. This essentially follows from the linearity of $(\mathcal{Z} \cdot)$. Let us first recall the definition of multiplication in the Giry monad. Given a (polish) space \mathcal{S} and a Borel set $O \subseteq \mathcal{S}$, the multiplication $\mu_{\mathcal{S}}$ is defined by the expected value of the evaluation function $\epsilon_O: \mathbb{D}\mathcal{S} \rightarrow \mathbb{R}$ where $\epsilon_O \cdot \delta = \delta \cdot O$ for every $\delta: \mathbb{D}\mathcal{S}$. We have, for every $\Delta: \mathbb{D}^2 \mathcal{S}$,

$$(\mu_{\mathcal{S}} \cdot \Delta) \cdot O = \int \epsilon_O d\Delta . \quad (10)$$

We simply write ϵ_s when O is the singleton $\{s\}$. Recall also that given a measurable function $f: \mathcal{R} \rightarrow \mathcal{S}$, the push-forward of f by \mathbb{D} is $\mathbb{D}f: \mathbb{D}\mathcal{R} \rightarrow \mathbb{D}\mathcal{S}$ where, for every Borel set $O \subseteq \mathcal{S}$ and measure $\delta: \mathbb{D}\mathcal{R}$, we have

$$\mathbb{D}f \cdot \delta \cdot O = \delta \cdot (f^{-1} \cdot O) = \delta \cdot \{r \mid f(r) \in O\} . \quad (11)$$

A very useful result links the constructor \mathbb{D} with integrals. Let $f: \mathcal{R} \rightarrow \mathcal{S}$ and $g: \mathcal{S} \rightarrow \mathbb{R}$ be measurable functions and $\delta: \mathbb{D}\mathcal{R}$ be a Borel measure. We know that $\mathbb{D}f \cdot \delta \in \mathbb{D}\mathcal{S}$ so we can integrate g wrt it. In fact, Giry has proven in [13, Sec. 3 p.70] that

$$\int g d\mathbb{D}f \cdot \delta = \int g \circ f d\delta . \quad (12)$$

²¹ Here, $(\mathbb{D}, \mu, [-])$ is the usual Giry monad, see [13, 20, 22].

We are now ready to prove that all diagrams in Fig. 6 commute. For Fig. 6-(c), let $\Delta: \mathbb{D}^2 \mathcal{X}^2$ be an arbitrary hyper and $(z, x'): \mathcal{Z} \times \mathcal{X}$. On the one hand,

$$\begin{aligned} & (\mu_{\mathcal{Z} \times \mathcal{X}} \circ \mathbb{D}(Z \cdot) \cdot \Delta)_{z, x'} \\ = & \int \epsilon_{(z, x')} d\mathbb{D}(Z \cdot) \cdot \Delta && \text{"Eqn. (10)"} \\ = & \int \epsilon_{(z, x')} \circ (Z \cdot) d\Delta && \text{"Eqn. (12)"} \end{aligned}$$

On the other hand, we have

$$\begin{aligned} & (Z \cdot (\mu_{\mathcal{X}^2} \cdot \Delta))_{z, x'} \\ = & \sum_x Z_{z, x} (\mu_{\mathcal{X}^2} \cdot \Delta)_{x, x'} && \text{"Def. matrix multiplication"} \\ = & \sum_x Z_{z, x} \int \epsilon_{(x, x')} d\Delta && \text{"Eqn. (10)"} \\ = & \int \sum_x Z_{z, x} \epsilon_{(x, x')} d\Delta && \text{"Integration is linear"} \end{aligned}$$

Thus, it remains to show that the integrands in the last lines of the reasoning above are the same. Let $\delta: \mathbb{D} \mathcal{X}^2$, we have

$$(\epsilon_{(z, x')} \circ (Z \cdot)) \cdot \delta = (Z \cdot \delta)_{z, x'} = \sum_x Z_{z, x} \delta_{x, x'} = \left(\sum_x Z_{z, x} \epsilon_{(x, x')} \right) \cdot \delta$$

Hence, Fig. 6-(c) commutes and its \mathbb{D} -lift gives the inner-average commutative diagram.

We now prove that Fig. 6-(b) (outer-average) commutes. Let $\Gamma: \mathbb{D}^3 \mathcal{X}^2$ and O be a Borel subset of $\mathbb{D}^2(\mathcal{Z} \times \mathcal{X})$. On the one hand, we have

$$\begin{aligned} & (\mu_{\mathbb{D}(\mathcal{Z} \times \mathcal{X})} \circ \mathbb{D}^2(Z \cdot)) \cdot \Gamma \cdot O \\ & (\mu_{\mathbb{D}(\mathcal{Z} \times \mathcal{X})} \cdot (\mathbb{D}^2(Z \cdot) \cdot \Gamma)) \cdot O && \text{"Def. (o)"} \\ = & \int \epsilon_O d\mathbb{D}^2(Z \cdot) \cdot \Gamma && \text{"Eqn. (10)"} \\ = & \int \epsilon_O \circ \mathbb{D}(Z \cdot) d\Gamma && \text{"Eqn. (12)"} \end{aligned}$$

On the other hand, we have

$$\begin{aligned} & (\mathbb{D}(Z \cdot) \circ \mu_{\mathbb{D} \mathcal{X}^2}) \cdot \Gamma \cdot O \\ = & \mathbb{D}(Z \cdot) \cdot (\mu_{\mathbb{D} \mathcal{X}^2} \cdot \Gamma) \cdot O && \text{"Def. (o)"} \\ = & (\mu_{\mathbb{D} \mathcal{X}^2} \cdot \Gamma) \cdot ((Z \cdot)^{-1} \cdot O) && \text{"Def. push-forward (11)"} \\ = & \int \epsilon_{(Z \cdot)^{-1} \cdot O} d\Gamma && \text{"Eqn. (10)"} \end{aligned}$$

As before, it remains to show that the integrands are equal. For every hyper $\Delta: \mathbb{D} \mathcal{X}^2$, we have

$$\begin{aligned} & (\epsilon_O \circ \mathbb{D}(Z \cdot)) \cdot \Delta \\ = & \epsilon_O \cdot (\mathbb{D}(Z \cdot) \cdot \Delta) && \text{"Def. (o)"} \\ = & \mathbb{D}(Z \cdot) \cdot \Delta \cdot O && \text{"Def. } \epsilon \text{"} \\ = & \Delta \cdot ((Z \cdot)^{-1} \cdot O) && \text{"Def. push-forward (11)"} \\ = & \epsilon_{(Z \cdot)^{-1} \cdot O} \cdot \Delta && \text{"Def. } \epsilon \text{"} \end{aligned}$$

This concludes that Fig. 6-(b) commutes.

Conversely, assume $h_1^{\mathcal{Z}} \sqsubseteq h_2^{\mathcal{Z}}$ for every extension \mathcal{Z} . For $\mathcal{Z} = \mathcal{X}$, we have $h_1^{\mathcal{X}} \sqsubseteq h_2^{\mathcal{X}}$. Let $\pi: \mathbb{D} \mathcal{X}$ and $\Pi: \mathbb{D} \mathcal{X}^2$ such that Π has π on the diagonal and 0 anywhere else. It follows from Equation (9) that $h_i \cdot \pi = h_i^{\mathcal{X}} \cdot \Pi$ and thus $h_1 \sqsubseteq h_2$. ◀

Proof that $\llbracket - \rrbracket$ preserves sequential composition [from §5.2]

► **Lemma 13.** *Let H^1, H^2 be HMM matrices. We have $\llbracket H^1; H^2 \rrbracket = \llbracket H^1 \rrbracket; \llbracket H^2 \rrbracket^{\mathcal{X}}$.*

Proof. Let $\pi: \mathbb{D}\mathcal{X}$ be a prior. On the one hand, let $y_{1,2}$ be observations associated to the inner δ in $\llbracket H^1; H^2 \rrbracket \cdot \pi$. That is, for pair $x, x' \in \mathcal{X}$ of initial and final value:

$$\delta_{x,x'} = \frac{\pi_x(H^1; H^2)_{x, (y_1, y_2), x'}}{N_{H^1; H^2, \pi, (y_1, y_2)}}$$

where the normalisation constant is defined as in §2.1:

$$N_{H, \pi, y} = \sum_{x, x'} \pi_x H_{x, y, x'}. \quad (13)$$

On the other hand, let us have a look at the inners of $\llbracket H^1 \rrbracket; \llbracket H^2 \rrbracket^{\mathcal{X}} \cdot \pi$. Let one such inner be $\delta: \mathbb{D}(\mathcal{X} \times \mathcal{X}')$. By Equation (9), we have $\llbracket H^2 \rrbracket^{\mathcal{X}} \cdot \gamma = \mathbb{D}(Z \cdot \llbracket H^2 \rrbracket \cdot \vec{\gamma})$ where the right-conditional and right marginal satisfy $Z \langle \vec{\gamma} = \gamma$. Therefore, there exists some inner $\gamma: \mathbb{D}(\mathcal{X} \times \mathcal{X}')$ of $\llbracket H^1 \rrbracket \cdot \pi$ such that $\delta = Z \cdot \rho$ for some inner $\rho: \mathbb{D}\mathcal{X} \times \mathcal{X}'$ of $\llbracket H^2 \rrbracket \cdot \vec{\gamma}$ and $Z \langle \vec{\gamma} = \gamma$. Thus

$$\gamma_{x,u} = \frac{\pi_x H^1_{x, y_1, u}}{N_{H^1, \pi, y_1}} \quad \text{and} \quad \rho_{u, x'} = \frac{\vec{\gamma}_u H^2_{u, y_2, x'}}{N_{H^2, \vec{\gamma}, y_2}}$$

for some observations $y_{1,2}$. Hence

$$\begin{aligned} & \delta_{x,x'} \\ = & (Z \cdot \rho)_{x,x'} && \text{“}\delta = Z \cdot \rho\text{”} \\ = & \sum_u Z_{x,u} \rho_{u,x'} && \text{“Def. } Z\text{.”} \\ = & \sum_u Z_{x,u} \frac{\vec{\gamma}_u H^2_{u, y_2, x'}}{N_{H^2, \vec{\gamma}, y_2}} && \text{“Def. of } \rho \text{ above”} \\ = & \sum_u \gamma_{x,u} \frac{H^2_{u, y_2, x'}}{N_{H^2, \vec{\gamma}, y_2}} && \text{“}Z \langle \vec{\gamma} = \gamma\text{”} \\ = & \sum_u \frac{\pi_x H^1_{x, y_1, u}}{N_{H^1, \pi, y_1}} \frac{H^2_{u, y_2, x'}}{N_{H^2, \vec{\gamma}, y_2}} && \text{“Def. of } \gamma \text{ above”} \\ = & \frac{\pi_x (H^1; H^2)_{x, (y_1, y_2), x'}}{N_{H^1, \pi, y_1} N_{H^2, \vec{\gamma}, y_2}} && \text{“Def. } H^1; H^2\text{”} \end{aligned}$$

It remains to prove that $N_{H^1; H^2, \pi, (y_1, y_2)} = N_{H^1, \pi, y_1} N_{H^2, \vec{\gamma}, y_2}$ (this essentially tells us that it does not matter whether we apply the normalisation process for each component in the sequential composition or we leave it until the end and just carry out normal HMM multiplication at the matrix level).

$$\begin{aligned} & N_{H^1, \pi, y_1} N_{H^2, \vec{\gamma}, y_2} \\ = & \sum_{u, x'} \vec{\gamma}_u H^2_{u, y_2, x'} N_{H^1, \pi, y_1} && \text{“Def. } N_{H^2, \vec{\gamma}, y_2}\text{”} \\ = & \sum_{u, x'} \sum_x \gamma_{x,u} H^2_{u, y_2, x'} N_{H^1, \pi, y_1} && \text{“}\vec{\gamma}_u = \sum_x \gamma_{x,u}\text{”} \\ = & \sum_{u, x'} \sum_x \frac{\pi_x H^1_{x, y_1, u}}{N_{H^1, \pi, y_1}} H^2_{u, y_2, x'} N_{H^1, \pi, y_1} && \text{“Substituting } \gamma_{x,u}\text{”} \\ = & \sum_{u, x'} \sum_x \pi_x H^1_{x, y_1, u} H^2_{u, y_2, x'} && \text{“Simplification with } N_{H^1, \pi, y_1} \text{ (}\dagger\text{)”} \\ = & \sum_{x, x'} \pi_x \sum_u H^1_{x, y_1, u} H^2_{u, y_2, x'} && \text{“Arith.”} \\ = & \sum_{x, x'} \pi_x (H^1; H^2)_{x, (y_1, y_2), x'} && \text{“Def. } (H^1; H^2)_{x, (y_1, y_2), x'}\text{”} \\ = & N_{(H^1; H^2), \pi, (y_1, y_2)} && \text{“Def. of } N_{(H^1; H^2), \pi, (y_1, y_2)}\text{”} \end{aligned}$$

(†) This step assumes the normalising constant N_{H^1, π, y_1} is not zero. If it is zero, then the equality we want to prove clearly holds.

Hence, $\llbracket H^1; H^2 \rrbracket \cdot \pi$ and $(\llbracket H^1 \rrbracket; \llbracket H^2 \rrbracket) \cdot \pi$ have the exact same inners associated to the same respective outer probabilities. ◀

Proof that sequential composition respects behavioural equivalence

[from §5.2]

► **Theorem 14.** *If $H \sqsubseteq H'$ and $K \sqsubseteq K'$ then $H;K \sqsubseteq H';K'$.*

Let us firstly prove a very important property of *HMMs*: they transform gain functions. We construct, for every $g: \mathbb{G}_{\mathcal{W}} \mathcal{X}^2$, a gain function $g^H: \mathbb{G}_{\mathcal{Y} \rightarrow \mathcal{W}} \mathcal{X}^2$ where

$$g^H \cdot \sigma.(x, x') = \sum_{y, u} g \cdot (\sigma.y) \cdot (x, u) H_{x', y, u}, \quad (14)$$

for every strategy $\sigma: \mathcal{Y} \rightarrow \mathcal{W}$ and $x, x_0 \in \mathcal{X}$. We have the following properties of g^H .

► **Lemma 15.** *Let H, K be *HMMs* and $\pi: \mathbb{D}\mathcal{X}$ be a prior. For every gain function $g: \mathbb{G}_{\mathcal{Y} \rightarrow \mathcal{W}} \mathcal{X}^2$, we have:*

1. $V_g[\pi \rangle H] = V_{g^H}[\pi]$,
2. $g^{H;K} = (g^H)^K$.

Proof. For 1., we have

$$\begin{aligned} & \sum_y \max_{w: \mathcal{W}} \sum_{x, x'} g \cdot w.(x, x') H_{x, y, x'} \pi_x && \text{“Def. } V_g[\pi \rangle H]\text{”} \\ = & \max_{\sigma: \mathcal{Y} \rightarrow \mathcal{W}} \sum_y \sum_{x, x'} g \cdot (\sigma.y) \cdot (x, x') H_{x, y, x'} \pi_x && \text{“Swap max and } \sum_y\text{”} \\ = & \max_{\sigma: \mathcal{Y} \rightarrow \mathcal{W}} \sum_x g^H \cdot \sigma.(x, x) \pi_x && \text{“Def. } g^H\text{”} \\ = & V_{g^H}[\pi] && \text{“Def. } V_g[\pi]\text{”} \end{aligned}$$

For 2., let $\sigma: \mathcal{Y} \rightarrow \mathcal{Y} \rightarrow \mathcal{W}$. We have

$$\begin{aligned} & (g^K)^H \cdot \sigma.(x, x') && \text{“Def. } g^H\text{”} \\ = & \sum_{y, u} g^K \cdot (\sigma.y) \cdot (x, u) H_{x', y, u} && \text{“Def. } g^K\text{”} \\ = & \sum_{y, u} \sum_{y', v} g \cdot (\sigma.y.y') \cdot (x, v) K_{u, y, v} H_{x', y, u} && \text{“Swap sums”} \\ = & \sum_y \sum_{y', v} g \cdot (\sigma.y.y') \cdot (x, v) \sum_u H_{x', y, u} K_{u, y, v} && \text{“Def. } H; K\text{”} \\ = & \sum_{y, y'} \sum_v g \cdot (\sigma.y.y') \cdot (x, v) (H; K)_{x', (y, y'), v} && \text{“Def. } g^{H;K}\text{”} \\ = & g^{H;K} \cdot \sigma.(x, x') \end{aligned}$$

as required. ◀

Now, we can prove Thm. 14.

Proof of Thm. 14. Let $H \sqsubseteq H'$, $K \sqsubseteq K'$, g be a gain function and $\pi: \mathbb{D}\mathcal{X}$ be a prior. On the one hand, assume $H \sqsubseteq H'$ and let us show that $H;K \sqsubseteq H';K$.

$$\begin{aligned} & V_g[\pi \rangle (H;K)] && \text{“Lem. 15.1”} \\ = & V_{g^{H;K}}[\pi] && \text{“Lem. 15.2”} \\ = & V_{(g^K)^H}[\pi] && \text{“Lem. 15.1”} \\ = & V_{g^K}[\pi \rangle H] && \text{“Lem. 15.1”} \\ \geq & V_{g^K}[\pi \rangle H'] && \text{“} H \sqsubseteq H'\text{”} \\ = & V_g[\pi \rangle (H';K)] && \text{“Lem. 15.1 and Lem. 15.2”} \end{aligned}$$

On the other hand, assume $K \sqsubseteq K'$ and let us show that $H';K \sqsubseteq H';K'$. For each observation y of H' , we consider the associated inner $\delta^y: \mathbb{D}\mathcal{X}^2$ and outer $p^y: [0, 1]$ where

$$(\delta^y)_{x, x'} = \frac{\pi_x H'_{x, y, x'}}{\mathbb{N}_{H', y, \pi}}$$

where $\mathbb{N}_{H', y, \pi} = \sum_{x, x'} \pi_x H'_{x, y, x'}$ is the usual column normalising constant. Then

$$\begin{aligned}
& V_g[\pi)(H';K)] \\
= & \sum_y p^y V_g[\delta_y)K] && \text{“Def. } \delta^y\text{'s”} \\
\geq & \sum_y p^y V_g[\delta_y)K'] && \text{“} K \sqsubseteq K'\text{”} \\
= & V_g[\pi)(H';K')] && \text{“Def. } \delta^y\text{'s”}
\end{aligned}$$

Hence, $H; K \sqsubseteq H'; K'$ follows by transitivity of \sqsubseteq . ◀

G Reduction to Bayes vulnerability via a collateral context [from §5.3]

In §5.3 we asked whether the refinement relation \sqsubseteq from §5.1 was too strong, and recalled that in [19] a construction was given that reduced refinement to Bayes vulnerability, at least for closed systems without collateral correlations: if $P \not\sqsubseteq Q$ then there was a context “post-compose with program R ” and a prior π such that the two hypers

$$\Delta^P = \llbracket P; R \rrbracket . \pi \quad \text{and} \quad \Delta^Q = \llbracket Q; R \rrbracket . \pi$$

had *Bayes* vulnerabilities in particular “the wrong way around”, i.e. that $V_{gid}(\Delta^P) < V_{gid}(\Delta^Q)$. That justified the failure of refinement, for if $P \sqsubseteq Q$ then we cannot, for any prior and R , have that $Q; R$ is *more* Bayes vulnerable than $P; R$.

That argument does not however work directly for refinements that are rejected because of *collateral* correlations: the post-composing R does not have access to the initial states of P, Q if those programs modify the state. A very direct (and somewhat brutal) work-around would be to re-use the construction of [19] by putting P, Q into a context that preserved the input by first copying it into Z say, then executing P, Q resp. and finally introducing an R as before but one that operated not on X directly but rather on the preserved copy held in Z . Because neither of P, Q assign to Z , it would still contain X ’s original value; but *leaks* in P, Q from X would be reflected in different hypers $\mathbb{D}^2(Z \times \mathcal{X})$ resulting from each. That is, informally we would have a collateral context containing declaration `var Z: \mathcal{X}` and the code

$$Z := X; \quad P; \quad R_Z \qquad \text{and} \qquad Z := X; \quad Q; \quad R_Z ,$$

where R_Z is the distinguishing R from [19] but operating on Z instead of X .

A more convincing approach however is to use the gain function g that distinguishes P, Q to make a “pre” collateral correlation Π^g instead of a “post” assignment R^g . It has been shown that the Π^g can be derived from the R^g ; but a more direct route is the following.

Assume the programs P, Q are indistinguishable wrt. the hypers on their final states (else we could simply appeal to [19]), but that they differ wrt. their information-flow effect on the initial state. Using `chn` from Lem. 3 in §3.2, take `chn.P` and `chn.Q`, the effective channels of P, Q resp. and let Δ^P, Δ^Q resp. be the hypers resulting from the action of those channels respectively on the prior π . By assumption $\Delta^P \not\sqsubseteq \Delta^Q$, and so there is some gain-function g with $V_g(\Delta^P) < V_g(\Delta^Q)$. It is shown in [27], but considering only channels (not *HMM*-programs), that there is indeed a collateral correlation $\Pi^g: \mathbb{D}(\mathcal{W} \times \mathcal{X})$ such that

$$V_{gid}[\Pi^g \cdot \text{chn}.P] < V_{gid}[\Pi^g \cdot \text{chn}.Q] ,$$

where $[\Pi^g \cdot \text{chn}.P]$ is the hyper resulting from the joint distribution in $\mathbb{D}(\mathcal{W} \times \mathcal{Y})$ formed by the matrix multiplication of the joint distribution Π^g by the stochastic `chn.P`.

Thus we have shown that if the g -vulnerabilities of `chn.P` and `chn.Q` wrt. some $g: \mathcal{W} \rightarrow \mathcal{X} \rightarrow \mathbb{R}$ and $\pi: \mathbb{D}\mathcal{X}$ mandate $P \not\sqsubseteq Q$, then the collateral *Bayes* vulnerabilities of P, Q are the wrong way around for some collateral correlation Π^g (whose right-marginal is of course π); general g -discrimination is again, as in [19], reduced to Bayes-vulnerability discrimination and compositionality.